Nuxeo Enterprise Platform - Version 5.1 and 5.2

The Reference guide

5.1 / 5.2

Copyright © 2000-2008, Nuxeo SAS.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2; with Invariant Section “Commercial Support”, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is available at the URL: http://www.gnu.org/copyleft/fdl.html

Table of Contents

I. Introduction 1
1. Preface 2
1.1. What this Book Covers 2
1.2. What this book doesn't cover 2
1.3. Target Audience 2
1.4. About Nuxeo 2
1.5. About Open Source 2
2. Introduction 3
2.1. Enterprise Content Management 3
2.1.1. Why ECM? 3
2.2. The Nuxeo ECM platform 3
2.3. Introduction FAQ 3
2.3.1. What are Nuxeo EP 5, Nuxeo EP and Nuxeo RCP? 3
2.4. Intended audience 3
2.5. What this book covers 3
3. Getting Started 5
3.1. Prerequisites 5
3.2. Setting up your Eclipse 5
3.2.1. Enabling Maven project loading 5
3.2.2. Enabling Subversion checkout 5
3.3. Learning from the project sample 6
3.3.1. Checkout project 6
3.3.2. Understanding sample code 7
3.3.3. Declaring book document type 7
3.3.4. Regulating book states 8
3.3.5. Displaying book documents 8
3.3.6. Making book documents indexable and searchable 10
3.3.7. Enabling drag&drop creation 11
3.3.8. Listening for events 11
3.4. Starting a new project 12
3.5. Using Documentation 12
3.6. Other IDEs: IntelliJ IDEA and NetBeans 12
3.6.1. IDEA 12
3.6.2. NetBeans 12
4. General Overview 14
4.1. Introduction 14
4.1.1. Architecture Goals 14
4.1.2. Main concepts and design 17
4.2. Nuxeo Runtime: the Nuxeo EP component model 19
4.2.1. The motivations for the runtime layer 19
4.2.2. Extensible component model 20
4.2.3. Flexible deployment system 23
4.2.4. Extension points and Nuxeo 5 24
4.3. Nuxeo EP layered architecture 25
4.3.1. Layers in Nuxeo EP 25
4.3.2. API and Packaging impacts 27
4.3.3. Illustration of the layered architecture 27
4.4. Core Layer overview 27
4.4.1. Features of Nuxeo Core 28
4.4.2. Nuxeo Core main modules 29
4.4.3. Schemas and document types 29
4.4.4. Life cycle associated to documents 30
4.4.5. Security model 31
4.4.6. Core events system 32
4.4.7. Query system 32
4.4.8. Versioning system 32
4.4.9. Repository and SPI Model 33

4.4.10. DocumentModel 33
4.4.11. Proxies 34
4.4.12. Core API 34
4.5. Service Layer overview 34
4.5.1. Role of services in Nuxeo EP architecture 34
4.5.2. Services implementation patterns 35
4.5.3. Platform API 36
4.5.4. Adapters 36
4.5.5. Some examples of Nuxeo EP services 37
4.6. Web presentation layer overview 37
4.6.1. Technology choices 37
4.6.2. Componentized web application 37
5. Schemas and Documents 40
5.1. Introduction 40
5.1.1. Concepts 40
5.2. Schemas 40
5.3. Core Document Types 41
5.4. ECM Document Types 41
5.4.1. Label and Icon 42
5.4.2. Default view 42
5.4.3. Layout 42
5.4.4. Containment rules 43
5.4.5. Summary 43
II. Platform Services 44
6. Exception Handling 45
6.1. Introduction 45
6.2. Extension Points 45
6.2.1. requestdump 45
6.2.2. listener 45
6.2.3. errorhandlers 45
7. Actions, Views, Navigation URLs and JSF tags 47
7.1. Introduction 47
7.2. Actions 47
7.2.1. Concepts 47
7.2.2. Manage actions 47
7.3. Views 50
7.3.1. UI Views 50
7.3.2. Manage views 50
7.4. Navigation URLs 51
7.4.1. Document view codec service 51
7.4.2. URL policy service 52
7.4.3. Additional configuration 53
7.4.4. URL JSF tags 53
7.5. Nuxeo Document Lists Manager 54
7.6. Nuxeo File Manager 55
7.7. Nuxeo JSF tags 55
8. Layouts 56
8.1. Introduction 56
8.1.1. Layouts 56
8.1.2. Widgets 56
8.1.3. Widget types 56
8.1.4. Modes 56
8.2. Manage layouts 56
8.2.1. Layout registration 57
8.2.2. Layout definition 57
8.2.3. Widget definition 58
8.2.4. EL expressions in layouts and widgets 59
8.3. Document layouts 59
8.4. Layout display 60
8.5. Standard widget types 60
8.5.1. text 60
8.5.2. int 60

8.5.3. secret 61
8.5.4. textarea 61
8.5.5. datetime 61
8.5.6. template 61
8.5.7. file 61
8.5.8. htmltext 61
8.5.9. selectOneDirectory 61
8.5.10. selectManyDirectory 61
8.5.11. list 61
8.5.12. checkbox 62
8.6. Custom templates 62
8.6.1. Custom layout template 62
8.6.2. Custom widget template 63
8.6.3. Builtin templates to handle complex properties 64
8.7. Custom widget types 65
8.8. Generic layout usage 65
9. Event Listeners and Scheduling 67
9.1. Introduction 67
9.2. Concepts 67
9.3. Adding an event listener 67
9.4. Upgrading an event listener 68
9.5. Adding an event 68
9.6. From CoreEvents to JMS Messages 68
9.7. Adding a JMS message listener 69
9.8. Scheduling 70
10. User Notification Service 71
10.1. Introduction 71
10.2. Notification concept 71
10.3. Notification channels 71
10.4. E-mail notifications 71
11. Indexing & Searching 73
11.1. Introduction 73
11.2. Configuration 73
11.2.1. Concepts 73
11.2.2. The `indexableDocType` extension point 73
11.2.3. The `resource` extension point 73
11.2.4. Field configuration 74
11.2.5. Text fields and analyzers 74
11.2.6. Boolean attributes 75
11.2.7. Schema resources and fields without configuration 75
11.2.8. Schema resources 75
11.2.9. Automatic fields configuration 75
11.3. Programmatic Searching 76
11.3.1. Fields and literals 76
11.3.2. `WHERE` statements 76
11.4. The Compass plugin 78
11.4.1. Configuring Compass 78
11.4.2. Global configuration 78
11.4.3. Mappings for Nuxeo 79
11.4.4. Text fields behavior 80
11.5. Building a search UI with QueryModel 81
12. Look and feel 82
12.1. Introduction 82
12.2. Principle 82
12.3. Mechanism 83
12.3.1. The Elements 83
12.3.2. The format 84
12.3.3. The negotiation 85
12.3.4. The engine 85
12.3.5. Resource management 86
12.3.6. Application 86
12.4. Customizing the theme 86

12.4.1. Modifying the current theme using `theme-default.xml` 87
12.4.2. Modifying the current theme 89
12.4.3. Adding a new theme and its pages 93
13. Authentication, Users & Groups Management 96
13.1. Introduction 96
13.2. Users and Groups configuration 96
13.2.1. Schemas 96
13.2.2. Directories 97
13.2.3. UserManager 97
13.2.4. User Management Interface 99
13.3. Authentication 99
13.3.1. Authentication Framework Overview 99
13.3.2. Pluggable JAAS Login Module 100
13.3.3. Pluggable Web Authentication Filter 101
14. Security Policy Service 109
14.1. Introduction 109
14.2. Architecture 109
14.3. Policy contributions 109
14.3.1. Core policy contribution 109
14.3.2. Search policy contribution 110
15. Workflow & jBPM 112
15.1. Workflow in Nuxeo 5.1 112
15.1.1. Deploying process definitions 112
15.2. Workflow from Nuxeo 5.2: the jBPM Service 113
15.2.1. Introduction 113
15.2.2. jBPM service configuration 113
15.2.3. Document management 115
15.2.4. Default processes 116
15.2.5. Nuxeo jBPM How-to 118
16. Document Versioning 120
16.1. Setting the version of a document 120
16.2. Modifying automatically the version of a document 120
16.3. Accessing document from previous version 121
16.4. The versioning service implementation 121
17. Audit Service 122
17.1. Introduction 122
17.2. Features 122
17.3. Architecture 122
17.4. Retrieving entries 122
17.5. Contributing the audit service 123
17.5.1. Recording new events types 123
17.5.2. Recording additional informations 123
18. Directories and Vocabularies 124
18.1. Introduction 124
18.2. Directory with a Relational Database Management System (SQL) server as backend 124
18.3. Directory with an LDAP server as backend 125
18.3.1. Server definition 125
18.3.2. Directory declaration 126
18.4. Handling references between directory entries 127
18.4.1. References defined by a many-to-many SQL table 127
18.4.2. Static reference as a dn-valued LDAP attribute 127
18.4.3. Dynamic reference as a ldapUrl-valued LDAP attribute 127
18.4.4. Defining inverse references 127
18.5. Combining multiple directories into a single virtual directory 127
18.5.1. Multi-directory sources 128
18.5.2. Sub-directories 128
18.6. The Directory API 128
18.7. Building custom option lists in forms with vocabularies 128
19. Mimetype detection 130
19.1. Introduction 130
19.2. MimetypeRegistry 130
19.3. Mimetype sniffing 131

19.4. Invoking the mimetype detection 133
20. Content Transformation 134
20.1. Introduction 134
20.2. Plugins module 134
20.2.1. Creating a plugin 134
20.2.2. Declaring a plugin module 134
20.2.3. Using a transform plugin 135
20.3. Available transforms 136
20.3.1. Document conversion 136
20.3.2. Pdfbox 138
20.3.3. OLE objects extraction 138
20.3.4. Office files merger 140
20.3.5. XSL Transformation 141
21. Nuxeo Conversion Service 142
21.1. Conversion Service vs Transformation Service 142
21.1.1. Motivations for this API changes 142
21.1.2. What has been improved 142
21.1.3. About compatibility 142
21.2. Using Conversion Service 143
21.2.1. About `BlobHolder` 143
21.2.2. built-in converters 143
21.2.3. Conversion Service API 144
21.2.4. Configuring converter service 145
21.2.5. Contributing converters 145
21.2.6. Converters based on external command line tools 146
22. Relations 147
22.1. Introduction 147
22.2. Concepts 147
22.3. Configuration 147
22.3.1. Graph instances 147
22.3.2. Resource adapters 148
22.4. Manage relations 148
22.5. Display relations 148
22.6. Architecture overview 149
23. Placeful Configuration 150
23.1. Introduction 150
23.2. Using Placeful Configuration 150
23.3. Contributing a placeful configuration 150
23.4. Available storage 152
23.4.1. In memory storage 152
23.4.2. Directory storage 152
23.5. Exemple of extension definition 152
24. Monitoring Nuxeo 154
24.1. Integrating Nuxeo monitoring in your management system 154
24.1.1. Inventory (nx:*,management=inventory) 154
24.1.2. Metric (nx:*,metric=*,management=metric) 154
24.1.3. Quality (nx:*,usecase=*,management=usecase) 154
24.2. Integrating management in nuxeo server 154
24.2.1. nuxeo-runtime-management 154
24.2.2. nuxeo-platform-management 154
24.2.3. nuxeo-webengine-management 155
24.3. Contributing management 155
24.3.1. Publishing 155
24.3.2. Providing shortcuts 156
24.3.3. Reporting quality 157
III. Core Services 158
25. Nuxeo Runtime 159
25.1. Overview 159
25.1.1. Main Goals 159
25.1.2. Main Features 159
25.2. What is OSGi? 159
25.3. OSGi Support 160

25.3.1. Supported Features 160
25.3.2. Unsupported Features 161
25.3.3. Planned Features 161
25.4. Component Model 161
25.4.1. What are components? 161
25.4.2. Main Features 162
25.4.3. Planned Features 162
25.4.4. Adapting Components 162
25.4.5. Flexible Model 162
25.4.6. Component Life Cycle 163
25.4.7. Component Extensibility 163
25.5. Supported Host Platforms 165
25.5.1. JBoss Integration 165
25.5.2. Eclipse Integration 166
25.6. Using Nuxeo Runtime 166
25.6.1. Creating Components 166
25.6.2. Using components 169
25.6.3. XML Component Descriptors 173
25.7. Integration tests for Nuxeo Runtime applications 174
25.7.1. The NXRuntimeTestCase base class 174
25.7.2. Frequent patterns 175
25.8. Detailed Architecture 175
25.9. References 176
26. Nuxeo Core Documentation 178
26.1. TODO: BS 178
26.2. Overview 178
26.2.1. Main goals 178
26.2.2. Nuxeo Core Components 178
26.3. Nuxeo Core Architecture 179
26.3.1. Model Layer (or Internal API) 179
26.3.2. Implementation Layer 179
26.3.3. Facade Layer (or Public API) 179
26.3.4. Deployment 179
26.3.5. Client Session 180
26.4. The Repository Model 180
26.4.1. Document and Schemas 180
26.4.2. Document Facets 181
26.4.3. Document Annotations 181
26.4.4. Document Access Control 182
26.4.5. Life Cycle 184
26.4.6. Query Engine 186
26.4.7. The Public API 186
26.4.8. Integration with Applications Servers 188
26.5. Extension Points 188
26.5.1. Session Factories 188
26.5.2. LifeCycle Managers 188
27. Nuxeo Core Import / Export API 190
27.1. Export Format 190
27.1.1. document.xml format 190
27.1.2. Inlining Blobs 191
27.2. Document Pipe 192
27.3. Document Reader 192
27.4. Document Writer 193
27.5. Document Transformer 194
27.6. API Examples 194
27.6.1. Exporting data from a Nuxeo repository to a Zip archive 195
27.6.2. Importing data from a Zip archive to a Nuxeo repository 195
27.6.3. Export a single document as an XML with blobs inlined. 196
28. Nuxeo Event Service 197
28.1. Nuxeo event model 197
28.1.1. Event 197
28.1.2. EventContext 197

28.1.3. EventListener 198
28.1.4. Transactions and Events 198
28.1.5. EventBundle 198
28.1.6. PostCommitEventListener 198
28.2. Using Events 199
28.2.1. Firing Events 199
28.2.2. Contributing an EventListener 199
28.2.3. Contributing a PostCommitEventListener 200
28.3. JMS and Nuxeo Events 201
28.3.1. JMS integration 201
28.3.2. Enabling JMS bridge 201
28.3.3. From 5.1 event model to 5.2 201
29. Experimental Topics 203
29.1. Introduction 203
29.2. Runtime Support for Scripting Languages 203
29.2.1. Introduction 203
29.2.2. Supported languages 203
29.2.3. Running a Script 204
IV. SOA, Web Services and various integration solutions 206
30. The Nuxeo Restlet API 207
30.1. Restlet Integration 207
30.1.1. Restlet types in Nuxeo 5 207
30.1.2. Restlet URL and parameters mapping 207
30.1.3. Contributing a new restlet 208
30.2. Nuxeo default restlets 208
30.2.1. Browse restlet 208
30.2.2. Export restlet 209
30.2.3. Lock restlet 209
30.2.4. Plugin upload restlet 210
30.3. Nuxeo RestPack 211
30.3.1. Installing the RestPack 211
30.3.2. Restlets included in the RestPack 211
30.4. Nuxeo WebEngine Restlets 215
31. Nuxeo HTTP client 216
31.1. HTTP Client Library 216
31.2. HTTP client authentication 216
32. Web services 218
32.1. Audit web service 218
32.2. Remoting web service 218
32.3. Indexing gateway service 218
32.4. Metro web services 218
33. Nuxeo JSR 168 Integration 220
33.1. Overview 220
33.2. Testing Nuxeo Portlets 220
33.2.1. Prerequisites 220
33.2.2. Generate a sample project with nuxeo-archetype-portlet archetype 221
33.2.3. Test the newly created portlet 222
33.3. Developping Nuxeo Portlets 222
33.3.1. NuxeoPortlet class 222
33.3.2. Project from nuxeo-archetype-portlet archetype 222
33.3.3. portlet.xml 223
33.3.4. Restlets 223
33.4. Available portlets 223
33.4.1. Nuxeo Search Portlet 223
34. Desktop integration tools 225
34.1. Drag and drop browser extensions 225
34.1.1. Server side import service: the FileManagerService 225
34.1.2. Microsoft Internet Explorer plugin 225
34.1.3. Mozilla Firefox plugin 225
34.2. Online document editing with LiveEdit 225
34.2.1. Functional overview 225
34.2.2. Functional use cases 226

34.2.3. Architectural overview 227
34.2.4. The Web Service component 230
34.2.5. More on editor launch 230
34.2.6. More on pre- and post-editing actions 231
34.3. Configuring LiveEdit links 231
34.3.1. Configuration policies 231
34.3.2. Changing the configuration policy 232
35. Nuxeo WebDAV interface 233
35.1. WebDAV clients 233
35.1.1. Path vs displayName 233
35.1.2. Filesystem resource vs Nuxeo DocumentModel artifact 233
35.1.3. MS Web Folder client 233
35.2. Fooling WebDAV clients 233
35.2.1. Available hacks 234
35.2.2. Configuring Nuxeo WebDAV connector for each client. 234
35.3. Nuxeo EP WebDAV implementation 234
35.3.1. Nuxeo EP WebDAV-specific features 235
35.3.2. Known limitations 235
35.4. Using the Nuxeo WebDAV connector 236
35.4.1. Installing the WebDAV connector 236
35.4.2. Connecting a client to Nuxeo WebDAV connector 236
36. Reporting: Eclipse BIRT Driver 237
36.1. Overview 237
36.2. How to use it 237
36.3. Tomcat integration HOWTO 242
37. Nuxeo Flex Connector 244
37.1. Overview 244
37.2. Development environment 244
37.3. Build and Deploy 244
37.3.1. Sample Overview 244
37.4. Dive In 245
37.4.1. Data Services Configuration 245
37.4.2. Your First Flex application 246
37.4.3. Granite DS configuration 247
V. Administration overview 248
38. OS requirements, existing and recommended configuration 249
38.1. Required software 249
38.2. Recommended configuration 249
38.2.1. Hardware configuration 249
38.2.2. Default configuration 249
38.2.3. For optimal performances 249
38.3. Known working configurations 250
38.3.1. OS 250
38.3.2. JVM 250
38.3.3. Storage backends 250
38.3.4. LDAP 251
39. SMTP Server configuration 252
40. RDBMS Storage and Database Configuration 253
40.1. Storages in Nuxeo EP 253
40.2. Installing the JDBC driver 253
40.3. Configuring Nuxeo Core Storage 253
40.3.1. Visible Content Store configuration 253
40.3.2. JCR backend configuration 254
40.3.3. Set up your RDBMS 257
40.3.4. Start Nuxeo EP 257
40.4. Configuring Storage for other Nuxeo Services 257
40.4.1. Configuring datasources 258
40.4.2. Relation service configuration 259
40.4.3. Compass search engine dialect configuration 260
40.5. Setting up a new repository configuration 260
40.5.1. Add the new repository configuration 260
40.5.2. Declare the new repository to the platform 261

41. LDAP Integration 263
41.1. For users/groups storage backend 263
42. OpenOffice.org server installation 264
42.1. Installation 264
42.1.1. Start server 264
42.1.2. Parameters 264
42.1.3. Installing an extension 265
42.1.4. Notes 265
42.2. Running OpenOffice as a Daemon 265
42.2.1. Nuxeo OOo Daemon 265
42.2.2. Configuring Nuxeo OOo Daemon 266
43. Run Nuxeo EP with a specific IP binding 267
44. Backup, restore and reset 268
44.1. Backup 268
44.2. Backup before an upgrade 269
44.3. Restore 269
44.4. Reset 269
45. The Nuxeo Shell 270
45.1. Overview 270
45.2. User Manual 270
45.2.1. Command Options 272
45.2.2. Commands 273
45.3. Troubleshooting 277
45.3.1. Check listened IP 277
45.3.2. Check connected server 277
45.3.3. Multi-machine case 277
45.4. Extending the shell 277
45.4.1. Registering New Custom Commands 278
45.4.2. Java Code for the new commands 278
45.4.3. Building the shell plugin 278
45.4.4. Deploying the shell plugin 279
VI. Core developer guide 280
46. Coding and Design Guidelines 281
46.1. Introduction 281
46.2. External Coding Standards 281
46.3. Some points that need attention 281
46.3.1. Java code formating 281
46.3.2. XML code formatting 282
46.3.3. Design 283
46.3.4. Unit tests 283
46.3.5. Security 283
46.3.6. Naming convention 284
46.3.7. Information hiding 284
46.3.8. Use modern Java features 284
46.3.9. Logging 284
46.3.10. Documentation: Comments and Javadoc 285
46.3.11. Deprecation 286
46.4. Methodology tips 287
46.4.1. Use the power of your IDE (and its plugins) 287
46.4.2. Refactor 288
46.5. Important references 288
47. Development Tools and Process 289
47.1. Code Quality with Eclipse Plugins 289
47.1.1. Using Checkstyle 289
47.1.2. Using TPTP 289
47.1.3. Using FindBugs 289
47.2. Profiling with NetBeans Profiler 289
47.3. NXPointDoc Documentation tool 289
47.3.1. Documenting a component 290
47.3.2. Creating the NxPointDoc site 292
47.3.3. Browsing NxPointDoc 292
47.4. Quality Assurance with continuous integration 292

47.4.1. Rules and means 293
47.4.2. Quality directives for Nuxeo developers 294
47.5. Release process 296
47.5.1. Overview 296
47.5.2. Continuous integration coverage 296
47.5.3. Help testing release candidates 297
48. Packaging Nuxeo EAR 298
48.1. Introduction 298
48.2. Basic project structure 298
48.3. The EAR module 299
48.3.1. Assembly descriptor 299
48.3.2. Add some resources 302
48.4. Improve usability 302
48.4.1. Thanks to Maven 302
48.4.2. Thanks to Ant 303
48.5. Recommended multi-machine packagings 306
48.5.1. Bi-machine: stateful/stateless 306
49. Release Management 311
49.1. Introduction 311
49.2. Let's release! 311
49.2.1. Remove all dependencies on SNAPSHOT versions 311
49.2.2. Checkout the code and clean your repository 312
49.2.3. Test the release 312
49.2.4. Perform the release 312
49.2.5. You're done 313
VII. Add-ons 314
50. Add-ons 315
50.1. Introduction 315
51. Nuxeo Annotation Service 316
51.1. Introduction 316
51.1.1. W3C Annotea 316
51.1.2. Logical architecture overview 316
51.1.3. XPointer integration and extension. 316
51.2. Annotation Service Core 317
51.2.1. Overview 317
51.2.2. Implementation 317
51.2.3. Storage 317
51.2.4. uriResolver 318
51.2.5. urlPatternFilter 318
51.2.6. metadata 318
51.2.7. permissionManager and permissionMapper 318
51.2.8. annotabilityManager 319
51.2.9. UID Generation 319
51.2.10. Event management 319
51.3. NXAS Repository Plugin 319
51.3.1. Overview 320
51.3.2. Default contribution 320
51.3.3. Extension Point 321
51.3.4. documentEventListener 321
51.3.5. securityManager 321
51.3.6. jcrLifecycleEventId and graphManagerEventListener 321
51.4. NXAS Web 321
51.4.1. Overview 321
51.4.2. Extension Point 321
51.4.3. Configuration 322
51.5. Annotation Service Facade 322
51.6. Annotation Service HTTP Gateway 322
51.6.1. Overview 322
51.7. References 323
52. Virtual Navigation 324
52.1. Virtual Navigation 324
52.1.1. Virtual Navigation presentation 324

52.1.2. Virtual Navigation configuration 324
53. Metadata Extraction Service 328
53.1. Introduction 328
53.2. Metadata extraction module 328
53.2.1. Defining a contribution for metadata extraction 328
53.2.2. Specifying input parameters 329
53.2.3. Chaining extractions 329
53.2.4. Creating a plugin for metadata extraction 329
53.2.5. Using a metadata extraction plugin 329
54. Unicity Service 331
54.1. How does it works ? 331
54.2. What you need: 331
54.3. Configuration 331
54.4. Snippets 331
55. Nuxeo Mail Service 333
55.1. Presentation 333
55.2. Basic Use 333
55.2.1. Configuration 333
55.3. Advanced Use 334
56. Imaging 335
56.1. Introduction 335
56.2. Imaging API 335
56.3. Imaging transform 335
56.4. Imaging web 335
56.5. Imaging core 335
57. Nuxeo Preview Addon 336
57.1. Overview 336
57.2. Installing the Preview Addon 336
57.2.1. Deploy the Addon 336
57.2.2. Configure Transformers 336
57.3. Extensions and Pluggability 337
57.3.1. Pluggable Adapters 337
57.3.2. Pluggable HTML Transformers 338
57.4. Previews and URLs 338
58. Nuxeo Tiling service Addon 340
58.1. Overview 340
58.2. How tiles are defined and computed 340
58.3. Installing the picture tiling Addon 341
58.3.1. Requirements 341
58.3.2. Deploy the addon 342
58.3.3. Configuration 342
58.4. Testing the tiling service 342
58.4.1. URL and restAPI 342
58.4.2. simple test client 344
59. Nuxeo DocumentLink Addon 345
59.1. Overview 345
59.2. How does it work 345
59.3. DocumentLink and Repository 346
59.4. Using DocumentLink 346
59.4.1. Using the adapter 346
59.4.2. Using the DocRepository 347
60. Nuxeo Search Center 348
60.1. Overview 348
60.2. Installing and testing the search center 348
60.2.1. Building from source 348
60.2.2. Deploying on a Nuxeo server 348
60.2.3. Using the Search Center 349
60.3. How does it work 349
60.3.1. The GWT client 349
60.3.2. The JAX-RS server 349
60.3.3. The SearchCenterService 349
60.4. How to customize the list of available filters 350

60.4.1. FilterWidgets 350
60.4.2. FilterSets 350
60.5. How to build new GWT filter widgets 351
60.5.1. Setting up a GWT eclipse development environment 351
60.5.2. Existing filter widgets 352
60.5.3. TODO: layout and extension points 353
VIII. Annexes 354
A. FAQs 355
A.1. Deployment and Operations 355
B. Detailed Development Software Installation Instructions 358
B.1. Installing Java 5 358
B.1.1. Using the Sun Java Development Kit (Windows and linux) 358
B.1.2. Using a package management systems (Linux) 358
B.1.3. Manual installation (Linux) 359
B.1.4. Setting up JAVA_HOME (Windows, Linux, Mac OS) 359
B.2. Installing Ant 359
B.3. Installing Maven 360
B.3.1. What is Maven? 360
B.3.2. Installing Maven 360
B.3.3. Generate a new project with the nuxeo-archetype-start archetype 361
B.4. Installing JBoss AS 361
B.4.1. JBoss AS listening ports customization 362
B.4.2. Affected JBoss services 364
B.5. Installing a Subversion client 365
B.5.1. Generic subversion clients with linux 365
B.5.2. Windows 366
B.6. Installing mercurial 366
B.6.1. Linux 366
B.6.2. Windows 366
B.6.3. Mac OS 366
B.6.4. Setting a username 366
B.6.5. Activating pre-integrated extensions 366
B.6.6. Using forest 367
B.7. Chapter Key Point 367
C. Commercial Support 368
C.1. About Us 368
C.2. Contact information 368
C.2.1. General 368
C.2.2. France 368
C.2.3. UK 368

Part I. Introduction

Chapter 1. Preface

1.1. What this Book Covers

The primary focus of this book it the presentation of Nuxeo EP, from the perspective of its configuration and its architecture.

1.2. What this book doesn't cover

This book is not an end-user manual for Nuxeo. If you are interested in such a document, we recommend that you get it from `http://doc.nuxeo.org/`.

1.3. Target Audience

As a reference book for the Nuxeo platform, this book has several intended audiences:

- System integrators, who need to understand how to configure and adapt the Nuxeo platform to their customers needs.
- Third-party application developers, who will typically need both to customize the behavior of some components and to write new components to extend the functionalities of the platform.
- System administrators, who need to understand how to configure the platform for the specific needs of their hosting environment.
- Core developers, who work on the platform and need a reference documentation.

We fully understand that these different people with different concerns, and we plan to create, thanks to the modular structure of this document, more targeted books for these different people.

Readers of this book are presumed familiar with the Java 5 language and with the Java EE and XML technologies.

1.4. About Nuxeo

Founded in 2000, Nuxeo SAS is part of the "second wave" of open source companies, and focuses on developing and supporting applications, instead of system software or development tools.

By entering the ECM field early in 2002, Nuxeo has established itself as the leader of open source ECM, with customers for critical projects in the Government, Energy and Finance sectors. Nuxeo currently has 40 employees, about half of them developers.

Nuxeo has customers and partners in Europe, in Northern and Southern America, in Africa and in India.

1.5. About Open Source

Simply put, Open source is a better way to develop and support software.

Nuxeo fully embraces the open source vision, by fostering collaboration with a community of external contributors.

Chapter 2. Introduction

This chapter will give you an overview of ECM and the motivation for using the Nuxeo platform in your next ECM project.

2.1. Enterprise Content Management

According to the AIIM, the Association for Information and Image Management, ECM is defined as "the technologies used to capture, manage, store, preserve, and deliver content and documents related to organizational processes".

2.1.1. Why ECM?

A March 2005 white paper on “The Hidden Costs of Information Work” by the IDC in the United States found that the average office employee spent approximately one day per week organizing and filing documents that they used. This can equate to a considerable amount of cost and inefficiency. A further twelve hours was spent on managing document approval, managing document routing and publishing to other channels. Nine hours was spent searching for documents.

2.2. The Nuxeo ECM platform

2.3. Introduction FAQ

2.3.1. What are Nuxeo EP 5, Nuxeo EP and Nuxeo RCP?

Nuxeo EP 5 is the 5th version of the open source platforms developed by Nuxeo (the four previous ones where known as "CPS").

Nuxeo EP or "Enterprise Platform" is the server part of Nuxeo EP 5. It is a Java EE application intended to run in a standard Java EE 5 application server like JBoss. It can be accessed by end users from a web browser, from office productivity suites like MS-Office or OpenOffice.org, or from rich client developed using the Nuxeo RCP technology (see below).

Nuxeo RCP or "Rich Client Platform" is a platform for building rich client applications, that usually connect to a Nuxeo EP server.

Nuxeo EP and Nuxeo RCP run on top of a common runtime, "Nuxeo Runtime", and share a common set of core components, called "Nuxeo Core".

2.4. Intended audience

2.5. What this book covers

Part 1: Introduction

Part 2:

Part 3:

Part 4:

Chapter 3. Getting Started

This document describes how to setup an Eclipse project environment for working with Nuxeo EP. If you encounter a problem don't forget you can share your experience on the ecm mailing list .

3.1. Prerequisites

We assume from now on that on your computer, you have the following ready-to-use environment:

- Java Development Kit (JDK) 1.5.x (Nuxeo EP is known to work now with Java 6 or OpenJDK, but this is currently not supported)
- Eclipse 3.4.0 (aka ganymede) for J2EE developper
- A Nuxeo EP installation from the last release (download the installation wizard). You can also download a nightly build installer from `http://www.nuxeo.org/static/snapshots/` , but keep in mind that this one may be broken, so don't erase your working setup with a nightly build: the result is not guaranteed.

Note: The Nuxeo EP installer also installs and sets up the right version of JBoss on your computer, so you don't need to install it by your own.

If some of these software are not already set up on your computer, please refer to the “Detailed Development Software Installation Instructions” appendix in the nuxeo book, where you will find useful help.

3.2. Setting up your Eclipse

You should configure your Eclipse for loading the plugins listed above.

3.2.1. Enabling Maven project loading

Nuxeo projects are built using the popular *Maven* build tool. Nuxeo packages are well described as maven *artifacts*. Nuxeo provides you a Maven repository that contains all you need for coding and assembling Nuxeo applications. This way, Nuxeo insures you that you're always running the correct version of software.

Maven Integration for Eclipse provides tight integration for Maven into the IDE. It especially provides wizards for creating new Maven projects and dependency management for build path based on Maven's pom file.

Here are the references for the Eclipse Update Manager:

```
Name: Update Site for Maven Integration
         URL: http://m2eclipse.sonatype.org/update
```

3.2.2. Enabling Subversion checkout

The Subversive plug-in gives you the ability to work with this CVS successor version control system from the Eclipse workbench. You can work with Subversion repositories in almost exactly the same way you can now work with CVS repositories using the CVS plug-in bundled in the standard Eclipse distribution.

Here are the references for the Eclipse Update Manager:

```
Name: Subversive Update Site is a part of
       URL: Ganymede Update Site. Look at Help > Software Updates... > Available Software > Ganymede > Collabor
```

```
Name: Subversive SVN Connectors
       URL: http://www.polarion.org/projects/subversive/download/eclipse/2.0/ganymede-site/
```

3.3. Learning from the project sample

Nuxeo provides you a *sample project* that figures out main Nuxeo EP components uses by introducing the book document type. You will see how :

- declaring the book document type;
- regulating book document states;
- making book documents indexable and searchable;
- displaying book documents;
- enabling drag and drop book creation mode;
- handling book document creation events.

3.3.1. Checkout project

Run eclipse and create the project by running the “checkout maven project using SCM” wizard. The project SCM url is at org.nuxeo.project.sample. It can be quite long as it downloads all available package libraries and sources.

From the Java perspective, in the package explorer view, on the left side of the IDE, you can now see the Nuxeo project sample structure and you're able to browse all the classes without errors.

3.3.1.1. Running a project

Now that you have created your project, you would want to see what your (not so tough) work produced?

3.3.1.1.1. Generating the plugin

You should first generate the jar artifact. From the package explorer view, select the `nuxeo-project-sample` folder and launch a maven package build using the contextual menu : *Run as* > *Maven2 package*. The project's target folder will contain the generated plugin jar file “nuxeo-project-sample-XXXX.jar”.

3.3.1.1.2. Deploying the plugin

Then, you should copy the jar file in the good Nuxeo directory. Create a `build.properties` file on the base of the `build.properties.sample` so that `jboss.dir` points to your jboss path. Note that when using the Nuxeo EP installer, the JBOSS home directory is the same as the Nuxeo EP home directory.

Launch the dialog from *Run* > *External tools* > *External tools configuration...*. Select the “Ant Build” tool and create a *new* configuration using the contextual menu. Select the file `build.xml` as buildfile and specify `copy` as arguments.

You're now able to launch the copy of the generated plugin jar file using the external tools menu.

3.3.1.1.3. Running Nuxeo EP

Defining a server configuration for the Nuxeo runtime enables you to start and stop Nuxeo directly from eclipse. This can be achieved by launching the new server wizard from the servers view. You should select a Jboss v4.0 server type and provide the Nuxeo EP home directory as server home. You shoud also configure a starting timeout value suitable for your computer. A typical value would be 60s in the nuxeo server properties view.

Open the internal web browser view and go to `http://localhost:8080/nuxeo/` . Log in nuxeo with Admin signature (Administrator/Administrator), if your project was correctly deployed, you should notice that the default font and logo have changed.

When developing, you will find it useful not to have to restart jboss when performing changes in xhtml pages. Since releases 5.1.5 and 5.2-M2, you can add the line `facelets.REFRESH_PERIOD=2` to the `nuxeo.properties` file in the `nuxeo.ear/config` folder: pages will be refreshed within 2 seconds.

For hot deploying web sample resources, you should use the ant target `web`.

3.3.2. Understanding sample code

Here is the project source layout:

```
nuxeo-project-sample
`-- src
|-- main
|   |-- java ❶
|   |   `-- org
|   |       `-- nuxeo
|   |           `-- project
|   |               `-- sample
|   `-- resources
|       |-- META-INF -- MANIFEST.MF ❷
|       |-- OSGI-INF ❸
|       |   `-- l10n
|       |-- directories
|       |-- nuxeo.war ❹
|       |   |-- icons
|       |   |-- img
|       |   `-- incl
|       |       `-- tabs
|       |-- schemas
|       |-- themes
|       `-- workflows
|
`-- test ❺
|-- java
|   `-- org
|       `-- nuxeo
|           `-- ecm
|               `-- sample
`-- resources
`-- META-INF
`-- OSGI-INF
```

❷ The main entry point is the `MANIFEST.MF` file. It defines variables like `Provide-Package`, `Require-Bundle`, `Nuxeo-Require` and `Nuxeo-Component` that drive the Nuxeo EP runtime. They define where to load component definitions with their associated classes and resources.

❸ Component definition files are by convention located into the `OSGI-INF` folder. Nuxeo EP component definition are mainly achieved by defining extension points.

Note

The project sample makes a component decomposition that is very atomic. At the opposite, the whole extension points could be defined in a single component definition. It's an architectural choice.

❶ Extension points may reference sample's specific classes. They are all being defined in the `src/main/java` folder.

❹ Web book dedicated web templates are defined in the `nuxeo.war` folder.Note that these resources will be written into the `nuxeo.war` folder at deployment time.

❺ Starting and stopping Nuxeo EP for testing features would be time consuming. Fixture code based on Junit is located into the `src/test` folder.

3.3.3. Declaring book document type

In the Nuxeo sample, a book is composed by the four fields `isbn`, `rating`, `publicationDate` and `literals`. These fields are well defined in the book schema. Here is the corresponding XML snipset extracted from the `books.xsd` file:

```
...
        <xs:element name="isbn" type="xs:string"/>
```

```
          <xs:element name="rating" type="xs:int"/>
          <xs:element name="publicationDate" type="xs:date"/>
          <xs:element name="keywords" type="bk:stringArray"/>
          ...
```

Book schema and type are both declared in the `types` service using extensions. Here is the corresponding XML snipset extracted from the core-types-contrib.xml file

```
...
          <extension target="org.nuxeo.ecm.core.schema.TypeService"
          point="schema">
          <schema name="book" src="schemas/book.xsd" prefix="bk">
          </extension>
          ...
          <extension target="org.nuxeo.ecm.core.schema.TypeService"
          point="doctype">
          <doctype name="Book" extends="Document">
          ...
          <schema name="book" >
          ...
          </doctype>
          </extension>
          ...
```

3.3.4. Regulating book states

Books are regulated by the `standard` life cycle (`project`, `approved`, `obsolete`, `deleted` states). This is achieved by declaring to the `lifecycle` service an extension that assign the `default` lifecycle to book documents. Here is the corresponding XML snipset extracted from the `lifecycle-contrib.xml` file

```
...
          <extension
          target="org.nuxeo.ecm.core.lifecycle.LifeCycleService" point="types">
          <types>
          <type name="Book">default</type>
          </types>
          </extension>
          ...
```

3.3.5. Displaying book documents

Books are figured out using a specific icon and label. This is achieved by declaring to the `types` service the associated resources. Here is the corresponding XML snipset extracted from the `ui-types-contrib.xml` file

```
...
          <extension target="org.nuxeo.ecm.core.schema.TypeService"
          point="type">
          ...
          <type id="Book" coretype="Book">
          ...
          <label>Book</label>
          <icon>/icons/book.png</icon>
          ...
          </type>
          ...
          </extension>
          ...
```

Book documents are displayed under folders or workspaces. Theses containment rules are declared to the `types` service. Here is the XML snipset extracted from the `ui-types-contrib.xml` that defines them.

```
  ...
          <extension target="org.nuxeo.ecm.platform.types.TypeService"
          point="types">
          ...
          <type id="Folder" coretype="Folder">
          <subtypes>
          <type>Book</type>
          </subtypes>
          </type>

          <type id="Workspace" coretype="Workspace">
          <subtypes>
          <type>Book</type>
          </subtypes>
```

```
        </type>
        ...
        </extension>
        ...
```

ISBN book field is displayed onto book's content tab. This is achieved by registering into the `web layout` service a new layout that includes the ISBN book number. Here is the XML snipset extracted from the `layouts-contrib.xml` :

```
...
        <extension target="org.nuxeo.ecm.platform.forms.layout.WebLayoutManager"
        point="layouts">
        ...
        <layout name="book">
        ...
        <widget name="isbn" type="text">
        <labels>
        <label mode="any">ISBN</label>
        </labels>
        <fields>
        <field>bk:isbn</field>
        </fields>
        </widget>
        ...
        </layout>
        ...
        </extension>
        ...
```

The `types` service is configured for displaying the book layout. Here is the XML snipset extracted from the `ui-types-contrib.xml` :

```
...
        <extension target="org.nuxeo.ecm.platform.types.TypeService"
        point="types">
        ...
        <type id="Book" coretype="Book">
        ...
        <layouts mode="any">
        <layout>heading</layout>
        <layout>book</layout>
        <layout>file</layout>
        </layouts>
        ...
        </type>
        ...
        </extension>
        ...
```

Specific book behaviours are available onto book and workspace documents. Each corresponding tab is declared to the `actions` service. Note that tab's templates are defined into the folder `nuxeo.war/inc/tabs`. Note also that these templates get their model using the `BookManagerBean` which implements book behaviours. Here is the XML snipset extracted from the `actions-contrib.xml` file :

```
...
        <extension target="org.nuxeo.ecm.platform.actions.ActionService" point="actions">
        ...
        <action id="tab_book_view"
        link="/incl/tabs/book_view.xhtml" enabled="true"
        label="Book" icon="/icons/file.gif" order="9">
        ...
        </action>
        ...
        <action id="tab_folder_books_list"
        link="/incl/tabs/folder_books_view.xhtml" enabled="true"
        label="Books" icon="/icons/file.gif" order="15">
        ...
        </action>
        ...
        </extension>
        ...
```

Book documents are subject to rights management and the corresponding tab should be displayed. ere is the XML snipset extracted from the `actions-contrib.xml` file :

```
...
        <extension target="org.nuxeo.ecm.platform.actions.ActionService" point="actions">
        ...
```

```
          <filter id="rights" append="true">
          <rule grant="true">
          <type>Book</type>
          </rule>
          </filter>
          ...
          </extension>
          ...
```

3.3.6. Making book documents indexable and searchable

Book documents are made indexable by declaring an `indexableDocType` extension point to the search service in the `search-contrib.xml` file. Here is the XML snipset extracted from the `search-contrib.xml` file

```
...
          <extension
          target="org.nuxeo.ecm.core.search.service.SearchServiceImpl"
          point="indexableDocType">
          <indexableDocType indexAllSchemas="true" name="Book" />
          </extension>
          ...
```

Books are searched by keywords using using the `query model` service. Here is the XML snipset extracted from the `querymodel-contrib.xml` file :

```
...
          <extension
          target="org.nuxeo.ecm.core.search.api.client.querymodel.QueryModelService"
          point="model">
          ...
          <queryModel name="BOOK">
          ...
          <pattern>
          SELECT * FROM Document WHERE ecm:primaryType = 'Book' AND
          dc:subjects = ?
          </pattern>
          ...
          </queryModel>
          ...
          </extension>
          ...
```

Query models are retrieved and invoked by the `BookResultsProvider` which provide the criterias from the session context. The `BookResultsProvider` is registered into the `results provider` service. Here is the XML snipset extracted from the `resultsprovider-contrib.xml` file :

```
...
          <extension
          target="org.nuxeo.ecm.webapp.pagination.ResultsProviderService"
          point="model">

          <resultsProvider name="BOOK" farm="bookResultsProvider" />

          </extension>
          ...
```

Finally, books are searched and listed in the book's tab using the results provider cache. Here is the XHTML snipset extracted from the `folder_books_view.xhtnml` :

```
...
          <nxu:methodResult name="provider"
          value="#{resultsProvidersCache.get('BOOK')}">

          <h:form>
          <ui:decorate template="/pagination/pageNavigationControls.xhtml">
          <ui:param name="provider" value="#{provider}" />
          </ui:decorate>
          </h:form>

          <h:dataTable var="bookinfo" value="#{provider.currentPage}"
          ...
          </h:dataTable>
          ...
```

3.3.7. Enabling drag&drop creation

Dragging an image onto a workspace will create a book document. The `BookFileManagerPlugin` request the file manager service for listening to `jpeg` and `gif` images upload events. At it will be notified, it creates book document models, synthetizing titles based on the uploaded filename, normalizing them by invoking the `BookTitleService`. Here is the XML snipset extracted from the `fileManager-contrib.xml`

```
...
          <extension
          target="org.nuxeo.ecm.platform.filemanager.service.FileManagerService"
          point="plugins">
          ...
          <plugin name="book_plugin"
          class="org.nuxeo.project.sample.BookFileManagerPlugin">
          <filter>image/gif</filter>
          <filter>image/jpeg</filter>
          </plugin>
          ...
          </extension>
          ...
```

The `BookTitleService` is defined as a service. The `BookTitleDescriptor` defines the configuration format suitable for this service. Here is the XML snipset extracted from the `booktitle-service-contrib.xml` file

```
...
          <implementation class="org.nuxeo.project.sample.BookTitleServiceImpl" />

          <service>
          <provide interface="org.nuxeo.project.sample.BookTitleService" />
          </service>

          <extension-point name="title">
          <object class="org.nuxeo.project.sample.BookTitleDescriptor" />
          </extension-point>
          ...
```

3.3.8. Listening for events

Empty book document models are filled by default. At the other hand, the keywords dictionary is automatically filled using the ISBN numbers and titles. This is achieved by registering the two event listeners `BookEventListener` and `BookISBNEventListener` into the `event` service for receiving document events. Here is the XML snipset extracted from the `event-contrib.xml` file :

```
...
          <extension
          target="org.nuxeo.ecm.core.listener.CoreEventListenerService"
          point="listener">

          <listener name="book_create"
          class="org.nuxeo.project.sample.BookEventListener">
          <eventId>emptyDocumentModelCreated</eventId>
          </listener>

          <listener name="book_update_isbn"
          class="org.nuxeo.project.sample.BookISBNEventListener"
          order="157">
          <eventId>documentCreated</eventId>
          <eventId>documentModified</eventId>
          </listene>

          </extension>
          ...
```

The keywords dictionary is configured in the `directories-contrib.xml` file:

```
          <extension target="org.nuxeo.ecm.directory.sql.SQLDirectoryFactory"
          point="directories">

          <directory name="book_keywords">
          <schema>vocabulary</schema>
          <idField>id</idField>

          <dataSource>java:/nxsqldirectory</dataSource>
          <table>book_keywords</table>
          <!-- one of "never", "on_missing_columns" or "always" -->
          <createTablePolicy>on_missing_columns</createTablePolicy>
```

```
          <dataFile>directories/book_keywords.csv</dataFile>
          </directory>

          </extension>
```

3.4. Starting a new project

The goal of the *nuxeo-archetype-start* template is to setup a Nuxeo EP plugin development environment. The template provides: a maven layout for sources, tests and dependencies, an Ant target for deployment. It also customizes the web application a litte bit.

Start the “new maven project” wizard from the menu *File > New > Maven project*. Enable your eclipse for using Nuxeo EP's by addind a remote archetypes catalog using the configure dialog. The URL to be used is as follow http://maven.nuxeo.org/nuxeo-release/. Select the archetype type *nuxeo-archetype-start* and the nuxeo's version you want your project based. Set the artifact parameters according to your organisation rules.

3.5. Using Documentation

- The *Nuxeo Book* is getting to be the most complete source of information around Nuxeo EP, both for beginners and advanced developer. It is a good start.
- The extension point documentation is also very useful: although you may find it rough, it is the best way to evaluate the Nuxeo extensibility potential, and one should always start with a quick look around all the extension points, to "think Nuxeo" before starting a new project, and not reinventing the wheel.
- The wiki: we try to reference all the documentation from the wiki welcome page, and you will find tricks, howtos, etc. If you want to have a writer account to help update the content, ask on the Nuxeo's mailing list.

3.6. Other IDEs: IntelliJ IDEA and NetBeans

3.6.1. IDEA

IntelliJ IDEA from Jetbrains is a very lovable IDE for Java that has many fans in the Java developers community. It is unfortunately not open source.

To start using IDEA for coding on the Nuxeo project, you just need to type **mvn idea:idea** from your top-level source directory: Maven will download all the dependencies, then generate the configuration files needed by IDEA and you will be able open `nuxeo-ecm.ipr` from IDEA.

Note

At the time of this writing, IDEA will report some (spurious) compilation failures and you won't be able to compile the whole application from IDEA. You can still use IDEA with the current configuration to write Java code ("with pleasure"), and use Ant and/or Maven to build/deploy the application.

3.6.2. NetBeans

NetBeans is an open source IDE from Sun.

If you're using NetBeans 6 or later, you will be able to download a Maven 2 support plugin from NetBean's "Update Center", and with it, start coding on Nuxeo EP very quickly.

Chapter 4. General Overview

4.1. Introduction

4.1.1. Architecture Goals

When we started building Nuxeo EP, we defined several goals to achieve. Because these goals have a structural impact on Nuxeo EP platform it is important to understand them: it helps understanding the logic behind the platform.

4.1.1.1. Flexible deployment on multiple targets

An ECM platform like Nuxeo EP can be used in a lot of different cases.

The deployment of Nuxeo EP must be adapted to all these different cases:

- Standard ECM web application

  This is the most standard use case. The web browser is used to navigate in content repositories.

  - All services are deployed on an Application server
  - In order to be easily hostable, the platform needs to be compatible with several application servers

- Complex edition or rich media manipulation

  In this case having a rich client that seamlessly communicates with other desktop applications and offers a rich GUI is more comfortable than a simple web browser.

  - Interface is deployed on a rich client on the user desktop
  - Services and storage are handled on the server

- Offline usage

  In some cases, it is useful to be able to manipulate and contribute content without needing a network connection.

  - GUI and some services (like a local repository) need to be deployed on the client side
  - Server hosts the collaborative services (like the workflow) and the central repository

- Distributed architecture

  In order to be able to address the needs of large decentralized organizations, Nuxeo EP must provide a way to be deployed on several servers on several locations

  - Scale out on several servers
  - Dedicate servers to some specific services
  - Have one unique Web application accessing several decentralized repositories

- Use Nuxeo EP components from another application

  When building a business application, it can be useful to integrate services from Nuxeo EP in order to address all content oriented needs of the application.

  - Provide web service API to access generic ECM services (including repository)

- Provide EJB3 remoting API to access generic ECM services (including repository)
- Provide POJO API to generic ECM services

There are certainly a lot of other use cases, but mainly the constraints are:

- Be able to choose the deployment platform: POJO vs Java EE

  As first deployment targets we choose

  - Eclipse RCP: a rich client solution that uses a POJO (OSGi) component model
  - JBoss Application Server: a Java EE 5 compliant application server

- Be able to choose the deployment location of each component: client side vs server side

  The idea is to be able to deploy a component on the server side or on the client side without having to change its code or its packaging

4.1.1.2. Leverage CPS experience

Before building Nuxeo EP we worked during several years on the Zope platform with the CPS solution. CPS was deployed for a lot different use cases and we learned a lot of good practices and design patterns. Even if Nuxeo EP is a full rewrite of our ECM platform, we want to keep as much as possible of CPS good concepts.

- Concept of schemas and documents

  Inside CPS most of the data manipulated was represented by a document object with a structure based on schemas.

  This concept is very interesting:

  - Schemas enforce structure constraints and data integrity but also permit some flexibility.

    When defining a schema you can specify what fields are compulsory, what are their data type, but you can also define some flexible part of the schema.

  - Share API and UI components for Documents, Users, Records ...

    Because the Document/Schema model is very flexible it can be used to manipulate different types of data: like Users, Records and standards documents.

    From the developer's perspective this permit using the same API and be able to reuse some UI components

    From the user's perspective it gives the application some consistency: because the same features and GUI can be used for all the data he manipulates.

- Actions and Views

  Because CPS was very pluggable, it was possible to easily define different views for each document type and also to let additional components contribute new actions or views on existing documents.

  Nuxeo EP has a similar concept of views and actions, even if technically speaking the technologies are different.

- Lazy fetching and caching

  Because ECM applications make a very intensive use of the repository and often need to fetch a lot of different documents to construct each page, the way the document retrieval is handled if very important to have a scalable application.

With CPS we worked a lot on caching and lazy fetching.

With Nuxeo EP we incorporated this requirement from the beginning:

- Distributed caching
- Lazy fetching on schemas and fields

4.1.1.3. Extensible platform based on components

CPS was constructed as a set of pluggable components relying on a common platform. This modularity has been maintained in the new platform. Deploying a new feature or component on the new platform is as simple as it was on the old one.

This requirement has a huge impact on the platform because the Java packaging model and the Java EE constraints are not directly compatible with it.

Adding a new component should be as simple as dropping a file or an archive in some directory without having to rebuild nor repackage the application.

This is important from the administrator point of view: be able to easily deploy new features.

This is also very important from the support point of view: be able to deploy customized components without taking the risk of forking the platform and maintain the possibility to upgrade the standards components.

4.1.1.4. Easily accessible development framework

The CPS framework was powerful but we know it was very complex to use. Not only because of the unusual CMF/Zope/Python programming model, but also because there was a lot of different concepts and you had to understand them all to be able to leverage the platform when building a new application on top of it.

Nuxeo EP aims at simplifying the task of the developer

- Clearly separate each layer

  The idea is to clearly separate presentation, processing and storage so that developers can concentrate on their task.

- Offer plugin API and SPI

  Nuxeo EP is constructed as a set of plugins so you can modify the behavior of the application by just contributing a new plugin. This is simpler because for common tasks we will offer a simple plugin API and the developer just has to implement the given interface without having to understand each part of the platform.

- Rely on JAVA standards

  We try to follow as much as possible all the Java standards when they are applicable. This will allow experienced Java developers to quickly contribute to the Nuxeo EP platform.

4.1.1.5. Leverage Java open source community

We know what it's like to have to build and maintain an entire framework starting from the application server. With the switch to the Java technology, we will use as much as possible existing open source components and focus on integrating them seamlessly in the ECM platform. Nuxeo EP is a complete integrated solution for building an ECM application, but Nuxeo won't write all infrastructure components. This approach will also make the platform more standards compliant.

Thus developers can optimize their Java/JEE and open source experience to use Nuxeo EP.

4.1.1.6. Make the platform ready for SI integration

Because ECM applications often need to be deeply integrated into the existing SI, Nuxeo EP will be easily integrable

- API for each reusable service or component

  Depending on the components, this API could be POJO, EJB3, or WebService, and in most cases it will be available in the three formats.

- Pluggable hooks into Nuxeo EP

  This mainly means synchronous or asynchronous events listener that are a great place to handle communication and synchronization between applications.

4.1.1.7. Future-proof design

The Nuxeo EP platform was rewritten from the ground with the switch to Java. But we don't plan to do this kind of work every couple of years, it won't be efficient neither for us, nor for the users of the platform. For that reason, we choose innovative Java technologies like OSGi, EJB3, JSF, Seam ....

4.1.2. Main concepts and design

All the design goals explained just before have a huge impact on the Nuxeo EP architecture. Before going into more details, here are the main concepts of Nuxeo EP architecture.

4.1.2.1. Layered architecture

Nuxeo EP is built of several layers, following at least the 3 tiers standard architecture

- Presentation layer

  Handles GUI interactions (in HTML, SWT ...)

- Service layer

  Service stack that offers all generic ECM services like workflow, relations, annotations, record management...

- Storage layer

  Handles all storage-oriented services like document storage, versioning, life cycle ....

Depending on the components, their complexity and the needed pluggability, there can be more that 3 layers.

This layering of all the components brings Nuxeo EP the following advantages

- Choose the deployment target for each part of a component

  By separating clearly the different parts of a feature, you can choose what part to deploy on the client and what part to deploy on a server.

- Clear API separation

  Each layer will provide its own API stack

- Components are easier to reuse

  Because the service and storage layers are not bound to a GUI, they are more generic and then more reusable

Thanks to this separation in component families you can easily extract from Nuxeo EP the components you need for your application.

If you need to include Document storage facilities into your application you can just use Nuxeo EP Core: It will offer you all the needed feature to store, version and retrieve documents (or any structured but flexible dataset). If you also need process management and workflow you can also use Nuxeo EP Workflow service. And finally, if you want to have a Web application to browse and manage your data, you can reuse the Nuxeo EP Web layer.

4.1.2.2. Deployment services

The targeted platform do not provide the same mechanism to handle all the deployment tasks:

- Packaging (Java EE vs OSGi)
- Dependency management
- Extension management

Because of these differences, Nuxeo EP provides a unified deployment service that hides the specificity of the target platform. This is also a way to add a pluggable component deployment system to some platform that don't handle this (like Java EE).

This is one of the motivation for the Nuxeo Runtime that will be quickly introduce later in this document.

4.1.2.3. Extensible component model

In Nuxeo EP, an ECM application is seen as an assembly of components.

This assembly will include:

- Existing generic Nuxeo EP Components
- Extensions or configurations contributing to generic Nuxeo EP components
- Specific components and configuration

Inside Nuxeo EP each feature is implemented by a one or several reusable components and services. A feature may be implemented completely at storage level, or may require a dedicated service and a dedicated GUI.

Nuxeo EP Web application is a default distribution of a set of ECM components. This can be used "as is" or can be the base for making a business ECM application.

- If you need to remove a feature

  Just remove the component or deploy a configuration for disabling it.

- If you need to change the default behavior of one component

  You can deploy a new configuration for the component .

  - Declare a new Schema or define a document type
  - Configure the versioning policy
  - Deploy new workflow
  - ...

  This configuration may use an extension point to contribute the new behavior.

- Contribute a new security policy
- Contribute a new event handler
- Deploy a new View on a document
- ...

- If you need to add a completely new feature you can make your own component.

  First check that there is no generic Nuxeo EP component available that could help you in your task (all components are not deployed in the default webapp).

4.1.2.4. Use of innovative Java EE technology

Here is a quick list of the Java technology we use inside Nuxeo EP platform:

- Java 5
- Java EE 5: JSF and EJB3
- OSGi component model
- A lot a innovative open source projects
  - JBoss Seam, Trinidad and Ajax4JSF on the web layer
  - jBPM for the default workflow engine implementation
  - Lucene for the default search engine implementation
  - Jackrabbit JSR-170 repository for the default storage back end implementation
  - JenaRDF for the relation framework
  - ...

4.2. Nuxeo Runtime: the Nuxeo EP component model

4.2.1. The motivations for the runtime layer

Building the Nuxeo Runtime was one of the first task we started. This is one of the main infrastructure component or Nuxeo EP architecture.

This paragraph will give you a quick overview of the Nuxeo Runtime, a more detailed technical presentation can be found in an other chapter of this book.

4.2.1.1. Host platform transparency

Because most of Nuxeo EP components are shared by Nuxeo RCP (OSGI/RCP) and Nuxeo EP (Java EE), an abstraction layer is required so the components can use transparently the components services independently from the underlying infrastructure.

Nuxeo Runtime provides an abstraction layer on top of the target host platform. Depending on the target host platform, this Runtime layer may be very thin.

Nuxeo Runtime already supports Equinox (Eclipse RCP OSGi layer) and JBoss 4.x (JMX). The port of Nuxeo Runtime to other Java EE application server is in progress, we already have a part of Nuxeo EP components

that can be deployed on top of SUN Glassfish application server. Technically speaking, the port of Nuxeo Runtime could be done on any JEE5 compliant platform and will be almost straightforward for any platform that supports natively the OSGi component model.

4.2.1.2. Overcome Java EE model limitations

Java EE is a great standard, but it was not designed for a component based framework: it is not modular at all.

- Java EE deployment model limitations
  - Most Java EE deployment descriptors are monolithic

    For example, the web.xml descriptor is a unique XML file. If you want to deploy an additional component that needs to declare a new Java module you are stuck. You have to make one version of the web.xml for your custom configuration. For Nuxeo EP platform, this constraint is not possible:

    - Components don't know each other

      Because there are a lot of optional components, we can't have a fixed configuration that fits all.

    - We can make a version of the web.xml for each possible distribution

      There are too many optional components to build one static web.xml for each possible combination.

    This problem with the web.xml is of course also true for a lot of standard descriptors (application.xml, faces-config.xml, persistence.xml, ejb-jar.xml ....)

  - One archive for one web application

    We have here the exact same problem than with the web.xml. additional components can contribute new web pages, new web components ... We can have a monolithic web archive.

  - No dependency declaration

    Inside Java EE there is no standard way to declare the dependency between components.

    Because Nuxeo EP is extensible and has a plugin model, we need that feature. A contribution is dependent on the component it contribute to:

    - Contribution is only activated if/when the target component is activated
    - The contribution must be deployed after the target component as it may override some configuration
- Java EE component model limitations
  - Unable to deploy a new component without rebuilding the whole package

    If you take a .ear archive and want to add a new component, you have to rebuild a new ear.

  - No support for versionned components

Nuxeo Runtime provides an extensible component model that supports all these feature. It also handles the deployment of these components on the target host platform.

4.2.2. Extensible component model

Nuxeo Runtime provides the component model for the platform.

This component model is heavily based on OSGi and provides the following features:

- Platform agnostic component model

Can be deployed on POJO and Java EE platforms

- Supports dependencies management

  Components explicitly declare their requirements and are deployed and activated by respecting the inferred dependency chain.

- Includes a plugin model

  To let you easily configure and contribute to deployed components

- A POJO test environment

  Nuxeo Runtime components can be unit tested using JUnit without the need of a specific container.

4.2.2.1. The OSGi component model

OSGi (Open Services Gateway initiative) is a great standard for components based Java architecture.

OSGi provides out of the box the following features:

- Dependencies declaration and management

  A component gets activated only when the needed requirements are fulfilled

- Modular deployment system
  - Manage bundles
  - Manage fragments (sub parts of a master bundle)
  - an OSGi bundle can define one or several services
- A system to identify and lookup for a component

For Nuxeo EP, OSGi standard provides a lot of the needed features. This is the reason why Nuxeo Runtime is based on OSGi, in fact Nuxeo Runtime component model is a subset of OSGi specification.

To ensure platform transparency, Nuxeo Runtime provides adapters for each target platform to help it support OSGi components.

- This adopter layer is very thin on Equinox (Eclipse RCP) since the underlying platform is already OSGi compliant.

- This adapter may be more complex for platform that are not aware of OSGi (JBoss 4.x or Glassfish)

  In this case, the runtime adapter will handle all OSGi logic and deploy the components as native platform components. For examples, on JBoss 4.x, Runtime components are deployed as JMX MBeans.

4.2.2.2. Extension points

OSGi does not define a plugin model, but the Eclipse implementation (Equinox) does provide an extension point system.

Because we used a lot the Eclipse Extension Point system and we liked it, Nuxeo Runtime also includes an Extension Point system.

Basically every Nuxeo Component can:

- declare its dependencies

The component will also be activated after all needed components

- declare exposed extension points

  Each components can define extension points that other components can use to contribute configuration or code.

- declare contribution to other components

These declarations are handled by the OSGi deployment descriptor (MANIFEST.MF)

```
Manifest-Version: 1.0
Bundle-ManifestVersion: 2
Bundle-Name: Nuxeo ECM Core
Bundle-SymbolicName: org.nuxeo.ecm.core;singleton:=true
Bundle-Version: 1.0.0
Bundle-Vendor: Nuxeo
Bundle-Localization: bundle
Bundle-Activator: org.nuxeo.ecm.core.NXCoreActivator
Bundle-ClassPath: ., lib/xsom.jar,
  lib/connector-api.jar,
  lib/java-cup-v11a.jar
Export-Package: org.nuxeo.ecm.core,
  org.nuxeo.ecm.core.api,
  org.nuxeo.ecm.core.api.local,
  org.nuxeo.ecm.core.jca,
  org.nuxeo.ecm.core.lifecycle,
  org.nuxeo.ecm.core.model
Require-Bundle: org.nuxeo.ecm.core.api,
  org.nuxeo.runtime
Nuxeo-Component: OSGI-INF/CoreService.xml,
  OSGI-INF/TypeService.xml,
  OSGI-INF/RepositoryService.xml,
  OSGI-INF/CoreExtensions.xml,
  OSGI-INF/SecurityService.xml
```

For example, this descriptor defines that

- this bundle depends on `org.nuxeo.ecm.core.api`
- this bundles contains Nuxeo components like `CoreServices.xml`

  The XML descriptor will be used to define new extension points or contribute to existing one.

An extension point is a way to declare that your component can be customized from the outside:

- Contribute configuration

  Activate or deactivate a component. Define resources for a given service.

- Contribute code and behavior

  Extension points also give you the possibility to register plugins

Extension points and contribution to extension points are defined using a XML descriptor that has to be referenced in the `MANIFEST.MF`.

Here is a simple descriptor example:

```
<component name="org.nuxeo.ecm.core.listener.CoreEventListenerService">
  <require>org.nuxeo.ecm.core.repository.RepositoryService</require>
  <implementation class="org.nuxeo.ecm.core.listener.impl.CoreEventListenerServiceImpl"/>

  <extension-point name="listener">
    <object class="org.nuxeo.ecm.core.listener.extensions.CoreEventListenerDescriptor"/>
  </extension-point>
  <extension target="org.nuxeo.ecm.core.listener.CoreEventListenerService" point="listener">
    <listener name="nxruntimelistener" class="org.nuxeo.ecm.core.listener.impl.NXRuntimeEventListener" />
  </extension>
  <extension target="org.nuxeo.ecm.core.listener.CoreEventListenerService" point="listener">
```

```
    <listener name="lifecyclelistener" class="org.nuxeo.ecm.core.lifecycle.impl.LifeCycleListener" />
  </extension>
</component>
```

- This fragment depends on the Repository Service

  This fragment won't be loaded until a Nuxeo Repository is setup

- This fragment declares an extension point named listener

  This extension point let register plugins that will be invoked when a core event occurs.

  This extension point use `CoreEventListenerDescriptor` for descriptor.

- This fragment registers two contributions to the listener extension point

The contributions have to follow the descriptor defined by the target Extension Point. The descriptor defines what tags can be used when contributing.

The descriptor is simply defined by a Java class that uses annotations to defines how the XML descriptor will be used to create an Object Descriptor instance to pass to the extension point registration.

4.2.3. Flexible deployment system

Nuxeo Runtime also provides deployment services to manage how components are deployed and contribute to each other

- Dependencies management

  The dependencies are declared in the `MANIFEST.MF` and can also be defined in XML descriptors that hold contributions.

  The Nuxeo Runtime orders the component deployment in order to be sure the dependencies are respected. Components that have unresolved dependencies are simply not deployed

- Extension point contributions

  XML descriptors are referenced in the `MANIFEST.MF`. These descriptors make contributions to existing extension points or declare new extension points.

- Each component has its own deployment-fragment

  The deployment fragment defines

  - Contribution to configuration files

    For example contribute a navigation rule to `faces-config.xml` or a module declaration to `web.xml`.

    Nuxeo Runtime let you declare template files (like `web.xml`, `persistence.xml`) and let other component contribute to these files.

  - Installation instructions

    Some resources contributions (like i18n files or web pages) need more complex installation instructions because they need archives and files manipulations. Nuxeo Runtime provide basic commands to define how your components should be deployed

Here is a simple example of a deployment-fragment.

```
<fragment>
  <extension target="application#MODULE">
    <module> <ejb>${bundle.fileName}</ejb> </module>
```

```
  </extension>
  <extension target="faces-config#VALIDATOR">
    <validator>
      <validator-id>dueDateValidator</validator-id>
      <validator-class>org.nuxeo.ecm.platform.workflow.web.ui.jsf.DueDateValidator</validator-class>
    </validator>
  </extension>
  <install>
    <!-- unzip the war template -->
    <unzip from="${bundle.fileName}" to="/">
      <include>nuxeo.war/**</include>
    </unzip>
    <!-- create a temp dir -->
    <!-- be sure no directory with that name exists -->
    <delete path="nxworkflow-client.tmp" />
    <mkdir path="nxworkflow-client.tmp" />
    <unzip from="${bundle.fileName}" to="nxworkflow-client.tmp">
      <include>OSGI-INF/l10n/**</include>
    </unzip>
    <append from="nxworkflow-client.tmp/OSGI-INF/l10n/messages.properties"
              to="nuxeo.war/WEB-INF/classes/messages.properties" addNewLine="true" />
    <append from="nxworkflow-client.tmp/OSGI-INF/l10n/messages_en.properties"
              to="nuxeo.war/WEB-INF/classes/messages_en.properties" addNewLine="true" />
    <append from="nxworkflow-client.tmp/OSGI-INF/l10n/messages_fr.properties"
              to="nuxeo.war/WEB-INF/classes/messages_fr.properties" addNewLine="true" />
    <append from="nxworkflow-client.tmp/OSGI-INF/l10n/messages_de.properties"
              to="nuxeo.war/WEB-INF/classes/messages_de.properties" addNewLine="true" />
    <append from="nxworkflow-client.tmp/OSGI-INF/l10n/messages_it.properties"
              to="nuxeo.war/WEB-INF/classes/messages_it.properties" addNewLine="true" />
    <delete path="nxworkflow-client.tmp" />
  </install>
</fragment>
```

4.2.4. Extension points and Nuxeo 5

4.2.4.1. Some examples of extension point usage

Inside Nuxeo 5, extension points are used each time a behavior or a component needs to be configurable or pluggable.

Here are some examples of extension points used inside the Nuxeo 5 platform.

- Schemas and document types

  Inside Nuxeo 5 a document structure is defined by a set of XSD schemas. Schemas and Document Types are defined using an extension point.

- Storage repository

  Nuxeo core stores documents according to their type but independently of the low level storage back-end. The default back-end is Jackrabbit JCR implementation. Nuxeo Core exposes an extension point to define the storage back-end. We are working on an other repository implementation that will be pure SQL based.

- Security Management

  Nuxeo include a security manager that checks access rights on each single operation. The extension point system allow to have different implementation of the security manager depending on the project requirements:

  - Enforce data integrity: store security descriptors directly inside the document
  - Performance: store security descriptors in an externalized index
  - Corporate security policy: implement a specific security manager that will enforce business rules

- Event handlers

  Nuxeo platform lets you define custom Event handler for a very large variety of events related to content or processes. The event handler extension mechanism gives a powerful way to add new behavior to an

existing application

- You can modify the behavior of the application without changing its code
- The development model is easy because you have a very simple Interface to implement and you can use Nuxeo Core API to manipulate the data
- Event handlers can be synchronous or asynchronous

Nuxeo 5 itself uses the Event Handler system for a lot of generic and configurable service

- Automatic versioning: create a new version when a document is modified according to business rules
- Meta-data automatic update: update contributor lists, last modification date ...
- Meta-data extraction / synchronization: extract Meta-Data from MS-Office file, Picture ....

4.2.4.2. Nuxeo 5 Platform development model

Nuxeo 5 development model is heavily based on the usage of extension points. When a project requires specific features we try as much of possible to include it as an extension of the existing framework rather than writing separated specific component. This means make existing services more generic and more configurable and implement the project specific needs as a configuration or a plugin of a generic component using Extension Points.

4.3. Nuxeo EP layered architecture

4.3.1. Layers in Nuxeo EP

Nuxeo EP components are separated in 3 main layers: Core / Service / UI

From the **logical point of view** each layer is a group of components that provide the same nature of service:

- Storage oriented services: Nuxeo Core

  Nuxeo core provides all the storage services for managing documents

  - Repository service
  - Versioning service
  - Security service
  - Lifecycle service
  - Records storage (directories)
  - ...

- Content and process oriented services: Nuxeo Platform

  Nuxeo provides a stack of generic services that handle documents and provide content and process management features. Depending on the project requirement only a part of the existing services can be deployed.

  Typical Nuxeo EP platform services are:

  - Workflow management service

    - Relation management service
    - Archive management service
    - Notification service
    - ...
- Presentation service: UI Layer

  The UI layer is responsible for providing presentation services like

    - Displaying a view of a document
    - Displaying available actions according to context
    - Managing page flow on a process driven operation

  These services can be very generic (like the action manager) but can also be directly tied to a type of client (like the View generation can be bound to JSF/facelets for the web implementation)

The layer organization can also be seen as a **deployment strategy**

Thanks to the Nuxeo Runtime remoting features it is very easy to split the components on several JVM. But splitting some services can have a very bad effect on the global system performance.

Because of that, all the storage oriented services are inside the core. All components that have extensive usage of the repository and need multiple synchronous interaction with it are located in the core. This is especially true for all synchronous event handlers.

The services layer can itself be split in multiple deployment unit on multiple JVMs.

On the UI side all the services are logically deployed inside the same JVM. At least each JVM must have the minimum set of services to handle user interaction for the given application.

The components are also grouped by layers according to their **dependencies**.

Core Modules can depend on Core Internal API.

Generic ECM services can depend on Core external API and can depend on external optional library (like jBPM, Jena, OpenOffice.org ...).

UI services can rely on a client side API (like Servlet API) and share a common state associated to the user session.

Layers are also organized according to **deployment target**.

The Core layer is a POJO Layer with an optional EJB facade. The core can be embed in a client application.

The services are mostly implemented as POJO services so that they can be used as an embedded library. But some of them can depend on typical Application Server infrastructure (like JMS or EJB).

Inside the UI Layer most service are dedicated to a target platform: web (JSF/Seam), Eclipse RCP or other.

Because the layer organization has several constraints, the implementation of a unique feature is spread across several layers.

Typically a lot of transversal services is split in several sub-components in each layer in order to comply to deployment constraint and also to provide better reusability. For example, the Audit service is made of 3 main

parts:

- Core Event <=> JMS bridge (Core Layer)

  Forwards core events to JMS Bridge according to configuration.

- JMS Listener and JPA Logger (Service Layer)

  Message driven bean that writes logs in DB via JPA.

- Audit View (UI Layer)

  Generates HTML fragment that displays all events that occurred on a document.

4.3.2. API and Packaging impacts

The layer organization can also be seen in the API.

4.3.2.1. Core API

Most of the components forming the core are exposed via the `DocumentManager` / `CoreSession` interface. The interfaces and dependencies needed to access the Core services are packaged in a API package: even if there are several Core component, you have only one dependency and API package.

The idea is that for accessing the core, you will only need to use the `DocumentManager` to manipulate `DocumentModel`s (the document object artifact). Some core services can be directly accessed via the `DocumentModel` (like the life cycle or security data).

4.3.2.2. Service Stack API

Each service exposes its own API and then has its own API package. Service related data (like process data, relation data) are not directly hosted by the DocumentModel object but can be associated to it via adapters and facets.

4.3.2.3. UI API

The web layer can be very specific to the target application. Nuxeo EP provides a default web application and a set of base classes, utility classes and pluggable services to handle web navigation inside the content repository.

4.3.2.4. Packaging

Most features are made of several Java project and generate several Maven 2 artifact.

Nuxeo packaging and deployment system (Nuxeo Runtime, Platform API, Maven ...) leverage this separation to help you distributing the needed deployment unit according to your target physical platform.

4.3.3. Illustration of the layered architecture

XXX TODO

4.4. Core Layer overview

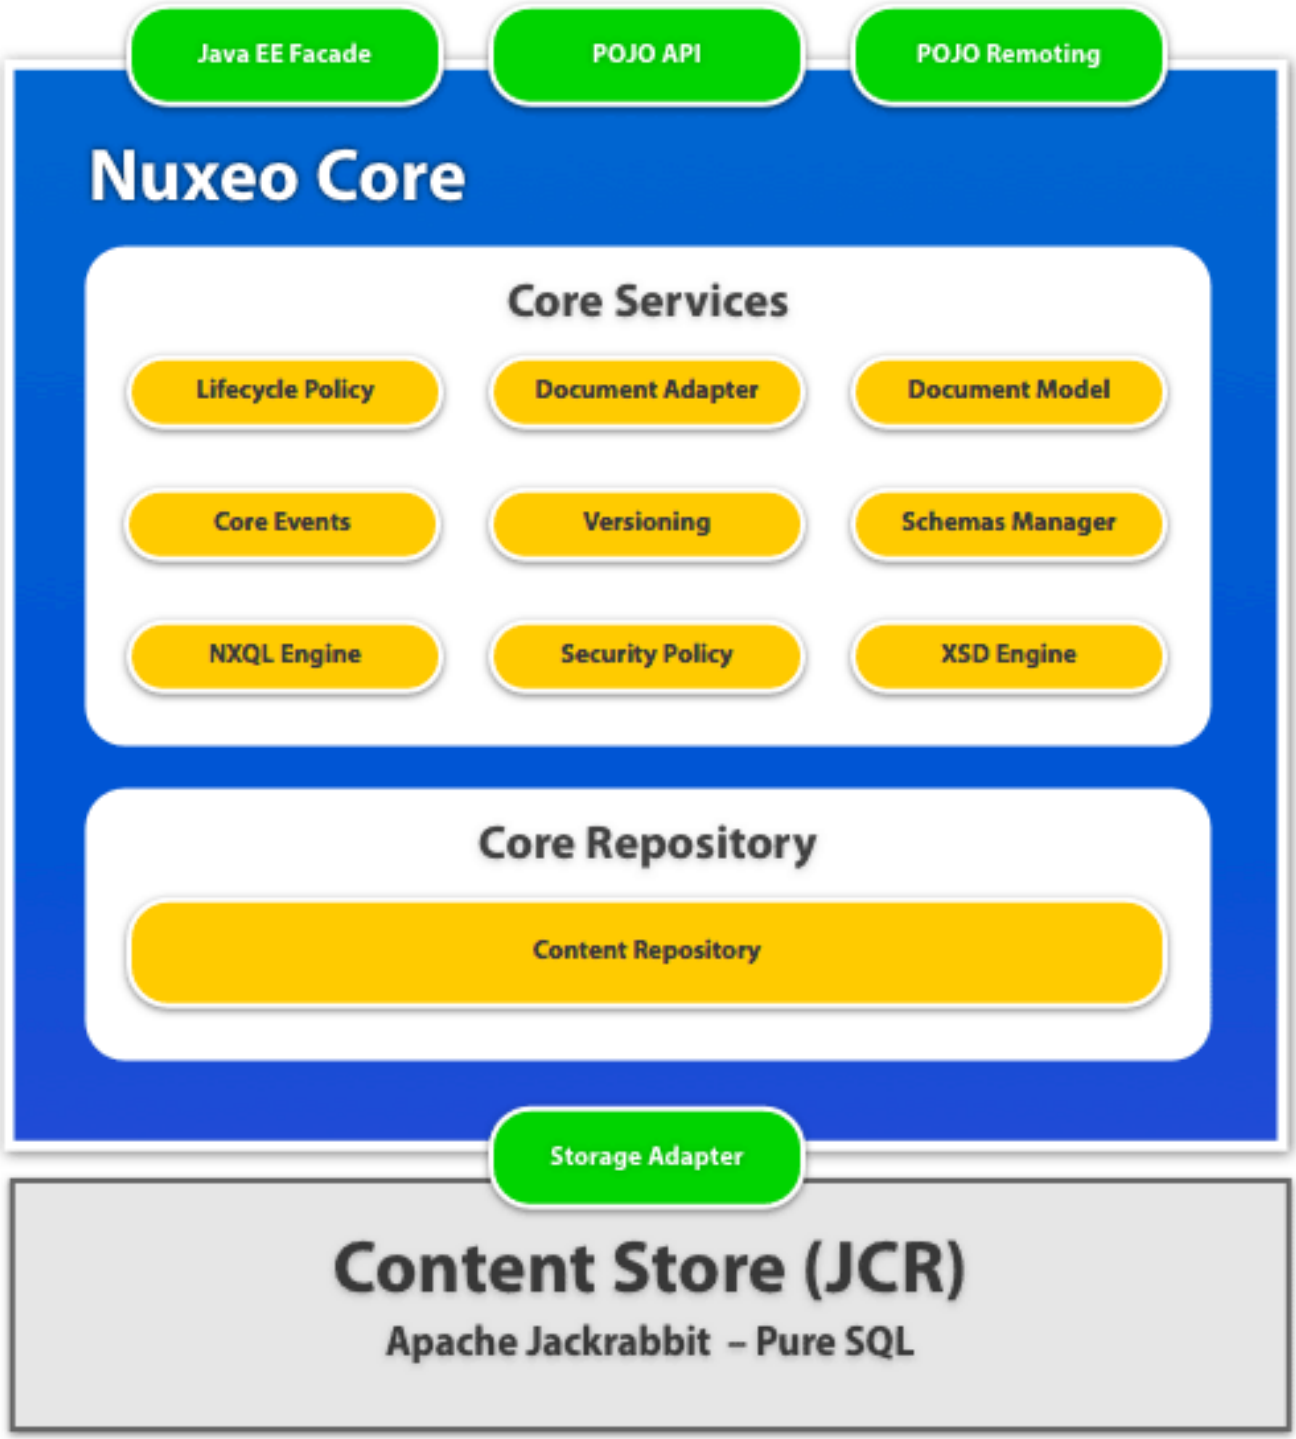

4.4.1. Features of Nuxeo Core

Nuxeo core provides all the storage services for managing documents.

- Schema service

  Lets you register XSD schemas and document types based on schemas.

- Repository service

  Lets you define one or more repository for storing your documents.

- Versioning service

Lets you configure how to store versions.

- Security service

  Manages data level security checks

- Lifecycle service

  Manages life cycle state of each document

4.4.2. Nuxeo Core main modules

4.4.2.1. Nuxeo Repository Service

The repository service lets you define new document repositories. Defining separated repositories for your documents is pretty much like defining separated Databases for your records.

Because Nuxeo Core defines a SPI on repository, you can configure how you want the repository to be implemented. For now, default implementation uses JSR-170 (Java Content Repository) reference implementation: Apache Jack Rabbit. In the future, we may provide other implementation of the Repository SPI (like native SQL DB or Object Oriented DB).

Even if for now there is only one Repository implementation available, using JCR implementation, you can configure how your data will be persisted: filesystem, xml or SQL Database. Please see "How to"s about repository configuration.

When defining a new repository, you can configure:

- The name.
- The configuration file

  For JCR, it lets you define persistence manager.

- The security manager

  Defines how security descriptors are stored in the repository (for now: org.nuxeo.ecm.core.repository.jcr.JCRSecurityManager)

- The repository factory

  Defines how the repository is created (for now: org.nuxeo.ecm.core.repository.jcr.JCRRepositoryFactory)

4.4.3. Schemas and document types

The repository enforces data integrity and consistency based on Document types definition.

Each document type is defined by:

- A name.
- An optional super document type (inheritance)
- A list of XSD schemas

  Defines storage structure

- A list of facets

  Simple markers used to define document behavior.

Here is a simple DocumentType declaration:

```
<extension target="org.nuxeo.ecm.core.schema.TypeService"
    point="doctype">
    <documentation>The core document types</documentation>
    <doctype name="Folder" extends="Document">
      <schema name="common" />
      <schema name="dublincore" />
      <facet name="Folderish" />
    </doctype>
    </extension>
```

For further explanation on Schemas and Document types, please see the dedicated section in this document.

4.4.4. Life cycle associated to documents

Inside Nuxeo repository each document may be associated with a life-cycle. The life-cycle defines the states a document may have and the possible transitions between these states. Here we are not talking about workflow or process, we just define the possible states of a document inside the system.

The Nuxeo Core contains a LifeCycleManager service that exposes several extension points:

- one for contribution Life-Cycle management engine

  (default one is called JCRLifeCycleManager and stores life-cycle related information directly inside the JSR 170 repository)

- one for contributing life-cycle definition

  This includes states and transitions. On 5.2, since 5.2.0, it is possible to define additional initial states to the default state, by adding a keyword to the state definition, for instance: <state name="approved" description="Content has been validated" initial="true">. The desired initial state can be passed in the document model context data used for creation: document.putContextData("initialLifecycleState", "approved").

```
<lifecycle name="default" lifecyclemanager="jcrlifecyclemanager"
  initial="project">

  <transitions>
    <transition name="approve" destinationState="approved">
      <description>Approve the content</description>
    </transition>
    <transition name="obsolete" destinationState="obsolete">
      <description>Content becomes obsolete</description>
    </transition>
    <transition name="delete" destinationState="deleted">
      <description>Move document to trash (temporary delete)</description>
    </transition>
    <transition name="undelete" destinationState="project">
      <description>Recover the document from trash</description>
    </transition>
    <transition name="backToProject" destinationState="project">
      <description>Recover the document from trash</description>
    </transition>
  </transitions>

  <states>
    <state name="project" description="Default state">
      <transitions>
        <transition>approve</transition>
        <transition>obsolete</transition>
        <transition>delete</transition>
      </transitions>
    </state>
    <state name="approved" description="Content has been validated">
      <transitions>
        <transition>delete</transition>
        <transition>backToProject</transition>
      </transitions>
    </state>
    <state name="obsolete" description="Content is obsolete">
      <transitions>
        <transition>delete</transition>
```

```
        <transition>backToProject</transition>
      </transitions>
    </state>
    <state name="deleted" description="Document is deleted">
      <transitions>
        <transition>undelete</transition>
      </transitions>
    </state>
  </states>

</lifecycle>
```

- one for binding life-cycle to document-types

  Here is an example

```
<lifecycle name="default" lifecyclemanager="jcrlifecyclemanager"
  initial="project">

  <transitions>
    <transition name="approve" destinationState="approved">
      <description>Approve the content</description>
    </transition>
    <transition name="obsolete" destinationState="obsolete">
      <description>Content becomes obsolete</description>
    </transition>
    <transition name="delete" destinationState="deleted">
      <description>Move document to trash (temporary delete)</description>
    </transition>
    <transition name="undelete" destinationState="project">
      <description>Recover the document from trash</description>
    </transition>
    <transition name="backToProject" destinationState="project">
      <description>Recover the document from trash</description>
    </transition>
  </transitions>

  <states>
    <state name="project" description="Default state">
      <transitions>
        <transition>approve</transition>
        <transition>obsolete</transition>
        <transition>delete</transition>
      </transitions>
    </state>
    <state name="approved" description="Content has been validated">
      <transitions>
        <transition>delete</transition>
        <transition>backToProject</transition>
      </transitions>
    </state>
    <state name="obsolete" description="Content is obsolete">
      <transitions>
        <transition>delete</transition>
        <transition>backToProject</transition>
      </transitions>
    </state>
    <state name="deleted" description="Document is deleted">
      <transitions>
        <transition>undelete</transition>
      </transitions>
    </state>
  </states>

</lifecycle>
```

Life-Cycle service is detailed later in this document.

4.4.5. Security model

Inside Nuxeo Repository security is always checked when accessing a document.

Nuxeo security model includes :

- Permissions

  (Read, Write, AddChildren, ...).

Permissions management is hierarchical (there are groups of permissions)

- ACE: Access Control Entry

  An ACE grants or denies a permission to a user or a group of users.

- ACL: Access Control List

  An ACL is a list of ACE.

- ACP: Access Control Policy

  An ACP is a stack of ACL. We use ACP because security can be bound to multiples rules: there can be a static ACL, an ACL that is driven by the workflow, and another one that is driven by business rules.

  Separating ACLs allows to easily reset the ACP when a process or a rules does not apply any more.

Inside the repository each single document can have an ACP. By default security descriptors are inherited from parent, but inheritance can be blocked when needed.

Security engine also lets you contribute custom policy services so that security management can include business rules.

Security model and policy service are described in details later in this document.

4.4.6. Core events system

When an event happens inside the repository (document creation, document modification, etc...), an event is sent to the event service that dispatches the notification to its listeners. Listeners can perform whatever action when receiving an event, this includes modifying the document on the fly.

As an example, part of the dublincore management logic is implemented as a CoreEvent listener: whenever a document is created or modified, creation date, modification date, author and contributors fields are automatically updated by a CoreEvent Listener.

Core Events system is explained in more details later in this document.

4.4.7. Query system

The Repository support a Query API to extract Documents using a SQL like query.

NXQL (the associated Query Language) is presented later in this document.

4.4.8. Versioning system

The documents in the repository can be versionned.

Nuxeo Core provides:

- A pluggable version storage manager

  This lets you define how versions and stored and what operations can be done on versions

- A pluggable versionning policy

  This lets you define rules and logic that drives when new versions must be created and how versions numbers are incremented.

The versionning system is explained in details later in this document.

4.4.9. Repository and SPI Model

Nuxeo Core exposes a repository API on top of Jackrabbit JSR170 repository.

Nuxeo repository is implemented using a SPI and extension point model: this basically means that a non JCR based repository plugin can be contributed. In fact, we have already started a native SQL repository implementation (that is not yet finished because we have no direct requirement for such a repository).

Nuxeo core can server several repository: it provides a extension point to declare additional repository: this means a single web application can use several document repository.

4.4.10. DocumentModel

Inside Nuxeo EP and especially inside the Core API the main data object is a Document.

Inside Nuxeo Core API, the object artifact used to represent a Document is called a DocumentModel.

The DocumentModel artifact encapsulates several useful features:

- Data Access over the network

  the DocumentModel encapsulate all access to Document internal fields, the DocumentModel can be sent over the network

- DocumentModel support lazy loading

  When fetched from the Core, a DocumentModel does not carries all document related information. Some data (called prefetch data) are always present, other data will be loaded (locally or remotely) from the core when needed.

  This feature is very important to reduce network and disk I/O when manipulating Document that contains a lot of big blob files (like video, music, images ...) .

- DocumentModel uses Core Streaming Service

  For files above 1 MB the DocumentModel uses the Core Streaming service.

- DocumentModel carries the security descriptors

  ACE/ACL/ACP are embedded inside the DocumentModel

- DocumentModels support an adapter service

  In addition of the data oriented interface, a DocumentModel can be associated with one or several Adapters that will expose a business oriented interface.

- DocumentModels embed lifecycle service access

- DocumentModels can have facets

  Facets are used to declare a behavior (Versionnable, HiddenInNavigation, Commentable...)

A DocumentModel can be located inside the repository using a DocumentRef. DocumentRef can be an IdRef (UUID in the case of the JCR Repository Implementation) or PathRef (absolute path).

DocumentModels also holds information about the Type of the Document and a set of flags to define some useful characteristics of the Document:

- isFolder

Defines if the targeted document is a container

- isVersion

  Defines if the targeted document is an historic version

- isProxy

  Defines if the targeted document is a Proxy (see below)

4.4.11. Proxies

A Proxy is a DocumentModel that points to a another one: very much like a symbolic link between 2 files.

Proxies are used when the same document must be accessible from several locations (paths) in the repository. This is typically the case when doing publishing: the same document can be visible in several sections. In order to avoid duplicating the data, we use proxies that point to same document.

A proxy can point to a checked out document (not yet associated to a version label) or to a versionned version (typical use-case of the publishing).

The proxy does not store document data: all data access are forwarded to the source document. But the Proxy does holds:

- its own security descriptors
- its own lifecycle information
- its own DocumentRef

4.4.12. Core API

4.5. Service Layer overview

4.5.1. Role of services in Nuxeo EP architecture

The service layer is an ECM services stack on top of the Nuxeo Repository. In a sense, the Repository itself is very much like any service of this layer, it just plays a central role.

This service layer is used for :

- adding new generic ECM services (Workflow, Relations, Audit ...)
- adding connectors to existing external services
- adding dedicated projects specific components when the requirements can not be integrated into a generic component

This service layer provides the API used by client applications (Webapp or RCP based application) to do their work.

This means that in this layer, services don't care about UI, navigation or pageflows: they simply explode an API to achieve document oriented tasks.

4.5.2. Services implementation patterns

Nuxeo platform provides a lot of different services, but they all fellow the same implementation pattern. This basically means that once you understand how works one service, you almost know how they all work.

As everything in the Nuxeo Platform the services use the Nuxeo Runtime component model.

A generic service will be composed of the following packages :

- A API package (usually called nuxeo-platform-XXX-api)

  Contains all interfaces and remotable objects.

  This package is required to be able to call the service from the same JVM or from a remote JVM.

- A POJO Runtime component (usually called nuxeo-platform-XXX-core)

  The Runtime component will implement the service business logic (ie: implement the API) and also expose Extensions Points.

  All the extensibility and pluggability is handled at runtime component level. This for example means that the API can be partially implemented via plugins.

- A EJB3 Facade (usually called nuxeo-platform-XXX-facade)

  The facade exposes the same API as the POJO component.

  The target of this facade is :

  - provide EJB3 remoting access to the API
  - integrate the service into JEE managed environment (JTA and JAAS)
  - leverage some additional features of the application server like JMS and Message Driven Bean
  - provide state management via Stateful Session Beans when needed

Typically, the POJO module will be a Nuxeo Runtime Component that inherit from DefaultComponent, provide extension points and implement a Service Interface.

```
public class RelationService extends DefaultComponent implements
        RelationManager { ...}
```

The deployment descriptor associated to the component will register the component, declare it as service provider and it may also declare extension points

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.platform.relations.services.RelationService">
  <implementation class="org.nuxeo.ecm.platform.relations.services.RelationService" />
  <service>
    <provide interface="org.nuxeo.ecm.platform.relations.api.RelationManager" />
  </service>
  <!-- declare here extension points -->
</component>
```

The facade will declare a EJB that implement the same service interface. In simple cases, the implementation simply delegates calls to the POJO component.

The facade package will also contain a contribution to the Runtime Service management to declare the service implementation.

```
<?xml version="1.0" encoding="UTF-8"?>
<component name="org.nuxeo.ecm.platform.relation.service.binding.contrib">
  <extension target="org.nuxeo.runtime.api.ServiceManagement" point="services">
      <documentation> Define the Relation bean as a platform service. </documentation>
      <service class="org.nuxeo.ecm.platform.relations.api.RelationManager" group="platform/relations">
```

```
        <locator>%RelationManagerBean</locator>
      </service>
    </extension>
</component>
```

Thanks to this declaration the POJO and the EJB Facade can now be used for providing the same interface based on a configuration of the framework and not on the client code.

This configuration is used when deploying Nuxeo components on several servers: platform configuration provides a way to define service groups and to bind them on physical servers.

4.5.3. Platform API

Each service provides its own API composed of a main interface and of the other interfaces and types that can be accessed.

The API package is unique for a given service, access to a remote EJB3 based service is the same as accessing the POJO service.

From the client point of view, accessing a service is very simple and independent from service location and implementation: this means not manual JNDI call. Everything is encapsulated in the Framework.getService runtime API.

```
RelationManager rm = Framework.getService(RelationManager.class);
```

The `framework.getService` will return the interface of the required service:

- This can be the POJO service (ie: Runtime Component based Service)
- This can be the local interface of the EJB3 service (using call by ref in JBoss)
- This can be the remote interface of the EJB3 service (using full network marshaling)

The choice of the implementation to return is left to the Nuxeo Runtime that will take the decision based on the platform configuration.

The client can explicitly ask for the POJO service via the Framework.getLocalService() API: this is typically used in the EJB Facade to delegate calls to the POJO implementation.

4.5.4. Adapters

DocumentModel adapters are a way to adapt the DocumentModel interface (that is purely data oriented) to a more business logic oriented interface.

In the pure Service logic, adding a comment to a document would look like this:

```
CommentManager cm = Framework.getService(CommentManager.class);
cm.createComment(doc,"my comment");
List<DocumentModel> comments = cm.getComments(doc);
```

DocumentModel adapter give the possibility to have a more object oriented API:

```
CommentableDoc cDoc = doc.getAdapter(CommentableDoc);
cDoc.addComment("my comment");
List<DocumentModel> comments = cDoc.getComments();
```

The difference may seem small, but documentModel adapter can be very handy

- to have a more clean and clear code

- to handle to caching at DocumentModel level
- to implement behavior and logic associating a Document and a Service

DocumenModelAdapters are declared using an extension point that defines the interface provided by the adapter and the factory class. DocumentModelAdapters can be associated to a Facet of the DocumentModel.

4.5.5. Some examples of Nuxeo EP services

4.6. Web presentation layer overview

4.6.1. Technology choices

4.6.1.1. Requirements

The requirements for the Nuxeo Web Framework are :

- A Powerful templating system that supports composition
- A modern MVC model that provides Widgets, Validators and Controllers
- A standard framework
- A set of Widgets libraries that allow reusing existing components
- Support for AJAX integration

4.6.1.2. The JSF/Facelets/Seam choice

Nuxeo Web Layer uses JSF (SUN RI 1.2) and Facelets as presentation layer: JSF is standard and very pluggable, Facelets is much more flexible and adapted to JSF than JSP.

NXThemes provides a flexible Theme and composition engine based on JSF and Facelets.

In the 5.1 version of the platform, Nuxeo Web Layer uses Apache Tomahawk and trinidad as components library and AJAX4JSF for Ajax integration. In the 5.2 version we will move to Rich Faces.

Nuxeo Web Layer also uses Seam Web Framework to handle all the ActionListeners.

Using Seam provides :

- Simplifications and helpers on JSF usage
- A context management framework
- Dependency injection
- Remoting to access ActionsListeners from JavaScript
- A built-in event system

4.6.2. Componentized web application

4.6.2.1. Requirements

Nuxeo Web Layer comes on top of a set of pluggable service.

Because this stack of services is very modular, so must be the web layer.

This basically mean that depending on the set of deployed services and on the configuration the web framework must provide a way

- to add, remove or customize views

  for example, if you don't need relations, you may want to remove the relations tab that is by default available on document

- to add or remove a action button or link

  the typical use case is removing actions that are bound to non deployed services or add new actions that are specific to your project

- to override an action listener

  you may want to change how some actions are handled by just overriding Nuxeo defaults

- to add or customize forms

  Adding fields or customizing forms used to display document is very useful

In order to fullfill these requirements, the key points of Nuxeo Web Framefulfill

- Context management to let components share some state
- Event system and dependency injection to let loosely coupled component collaborate
- A deployment system to let several components make one unique WebApp
- A set of pluggable services to configure the web application

4.6.2.2. Context management

Inside the web framework, each component will need to know at least

- what is the current navigation context

  This includes current document, current container, current Workspace, current Domain.

  This information is necessary because most of the service will display a view on the current document, and can fetch some configuration from the content hierarchy.

- who is the current user

  This includes identity and roles, but also its preferences and the set of documents he choose to work on

In some cases, this context information may be huge, and it's time consuming to recompute all this information at each request.

Inside Nuxeo Web Framework, Seam context management is used to store these data. Depending on the lifecycle of the data Session, Conversation or Event context are used.

4.6.2.3. Loosely coupled component

At some point the components of the web framework need to interact with each other. But because components

can be present or not depending on the deployment scenario, they can't call each other directly.

For that matter, the Nuxeo Web Framework uses a lot of Seam features:

- Seam's context is used to share some state between the components
- Seam's event system is used to let components notify each other
- Seam's dependency injection and Factory system is used to let component pull some data from each other without having to know each other

In order to facilitate Nuxeo Services integration into the web framework, we use the Seam Unwrap pattern to wrap Nuxeo Service into Seam Components that can be injected and manipulated as a standard Seam component.

4.6.2.4. Deployment

The Web Layer is composed of several components.

The main components are webapp-core (default webapp base) and ui-web (web framework). On top of these base components dedicated web components are deployed for each specific service.

For example, the workflow-service has its own web components package, so do relation-service, audit-service, comment-service and so on.

Each package contains a set of views, actions, and ActionsListeners that are dedicated to one service and integrate this service into the base webapp.

Because JEE standards require the webapp to be mono-bloc, we use the Nuxeo Runtime deployment service to assemble the target webapp at deployment time.

This deployment framework let you: override resources, contribute XML descriptors like `web.xml` from several components and manage deployment order.

4.6.2.5. Key web framework services

Chapter 5. Schemas and Documents

5.1. Introduction

This chapter presents the concepts of schemas and document types, which are used to define new documents in Nuxeo EP 5.

5.1.1. Concepts

In Nuxeo EP 5, a fundamental entity is the *document*. A file, a note, a vacation request, an expense report, but also a folder, a forum, can all be thought of as documents. Objects that contain documents, like a folder or a workspace, are also themselves documents.

Any given document has a *document type*. The document type is specified at creation time, and does not change during the lifetime of the document. When referring to the document type, a short string is often used, for instance “Note” or “Folder”.

A document type is the grouping of several *schemas*. A schema represents the names and structure (types) of a set of fields in a document. For instance, a commonly-used schema is the Dublin Core schema, which specifies a standard set of fields used for document metadata like the title, description, modification date, etc.

To create a new document type, we start by creating one ore more schemas that the document type will use.

The schema is defined in a `.xsd` file (which obeys the standard XML Schema syntax) and is registered by a contribution to a `schema` extension point.

The document type is registered through a contribution to another `doctype` extension point, and will specify which schemas it uses. We also register the document type as a high-level type used by the EJB and rendering layers.
“Core” document types and “ECM” component types should not be confused. The former live in the core Nuxeo EP packages, the latter belong to the high level components. Apart from their definition, most of the uses of "document types" refer to the “ECM” kind.

5.2. Schemas

A schema describes the names and types of some fields. The name is a simple string, like "title", and the type describes what kind of information it stores, like a string, an integer or a date.

First, we create a schema in the `resources/schemas/sample.xsd` file:

```
<?xml version="1.0"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
  targetNamespace="http://project.nuxeo.org/sample/schemas/sample/">
  <xs:element name="sample1" type="xs:string"/>
  <xs:element name="sample2" type="xs:string"/>
</xs:schema>
```

This schema defines two things:

- an XML namespace that will be associated with the schema,
- two elements and their type.

The two elements are *`sample1`* and *`sample2`*. They are both of type `xs:string`, which is a standard type defined by the XML Schema specification.

This schema has to be referenced by a configuration file so that the system knows it has to be read. The configuration file will be placed in `OSGI-INF/core-types-contrib.xml` (the name is just a convention):

```
<?xml version="1.0"?>
<component name="org.nuxeo.project.sample.coreTypes">
  <extension target="org.nuxeo.ecm.core.schema.TypeService" point="schema">
    <schema name="sample" src="schemas/sample.xsd"/>
  </extension>
</component>
```

We name our schema sample, and the `.xsd` schema was placed under `resources/` which means that at runtime, it will be available directly from the Jar, therefore we reference it through the path `schemas/sample.xsd`. The schema is registered through an extension point of the Nuxeo component `org.nuxeo.ecm.core.schema.TypeService` named `schema`. Our own extension component is given a name, `org.nuxeo.project.sample.coreTypes`, which is not very important as we only contribute to existing extension points and don't define new ones.

Finally, we tell the system that the `OSGI-INF/core-types-contrib.xml` file has to be read, by mentioning it in the Nuxeo-Component part of the `META-INF/MANIFEST.MF` file of our bundle:

```
Manifest-Version: 1.0
Bundle-ManifestVersion: 1
Bundle-Name: NXSample project
Bundle-SymbolicName: org.nuxeo.project.sample;singleton:=true
Bundle-Version: 1.0.0
Bundle-Vendor: Nuxeo
Require-Bundle: org.nuxeo.runtime,
 org.nuxeo.ecm.core.api,
 org.nuxeo.ecm.core
Provide-Package: org.nuxeo.project.sample
Nuxeo-Component: OSGI-INF/core-types-contrib.xml
```

5.3. Core Document Types

By itself, the schemas is not very useful. We associate it with a new “core” document type that we will be creating. In the same `OSGI-INF/core-types-contrib.xml` as above, we add the following:

```
<?xml version="1.0"?>
<component name="org.nuxeo.project.sample.coreTypes">
  <extension target="org.nuxeo.ecm.core.schema.TypeService" point="schema">
    <schema name="sample" src="schemas/sample.xsd" prefix="smp"/>
  </extension>
  <extension target="org.nuxeo.ecm.core.schema.TypeService" point="doctype">
    <doctype name="Sample" extends="Document">
      <schema name="common"/>
      <schema name="dublincore"/>
      <schema name="sample"/>
    </doctype>
  </extension>
</component>
```

The document type is defined through a contribution to an other extension point, *doctypes*, of the same extension component as before, `org.nuxeo.ecm.core.schema.TypeService`. We specify that our document type, `Sample`, will be an extension of the standard system type `Document`, and that it will be composed of three schemas, two standard ones and our specific one.

The Dublin Core schema that we use already contains standard metadata fields, like a title, a description, the modification date, the document contributors, etc. Adding it to our document type ensures that a minimal level of functionality will be present.

5.4. ECM Document Types

After the “core” document type, we need to create the “ECM” document type. This is done through a contribution to another extension point, that we place in the `OSGI-INF/ecm-types-contrib.xml`, the basic structure of this file is:

```
<?xml version="1.0"?>
<component name="org.nuxeo.project.sample.ecm.types">
  <extension target="org.nuxeo.ecm.platform.types.TypeService" point="types">
    <type id="Sample" coretype="Sample">
      ...
```

```
    </type>
  </extension>
</component>
```

As of this writing, the document type id has to be the same as the underlying core type; this restriction may be lifted in the future. The type element will contain all the information for this type, this is described below.

This extension file is added to `META-INF/MANIFEST.MF` so that it will be taken into account by the deployment mechanism:

```
Nuxeo-Component: OSGI-INF/core-types-contrib.xml,
 OSGI-INF/ecm-types-contrib.xml
```

Inside the type element, we provide various information, described below.

5.4.1. Label and Icon

```
<label>Sample document</label>
<icon>/icons/file.gif</icon>
```

The label and icon are used by the user interface, for instance in the creation page when a list of possible types is displayed. The icon is also used whenever a list of document wants to associate an icon with each.

5.4.2. Default view

```
<default-view>view_documents</default-view>
```

The default view specifies the name of the Facelet to use for this document if nothing is specified. This corresponds to a file that lives in the webapp, in this case `view_documents.xhtml` is a standard view defined in the base Nuxeo bundle, that takes care of displaying available tabs, and the document body according to the currently selected type. Changing it is not advised unless extremely nonstandard rendering is needed.

5.4.3. Layout

5.4.3.1. Configuration

Here is the new layout configuration:

```
<layouts mode="any">
  <layout>heading</layout>
  <layout>note</layout>
</layouts>
```

Layouts are defined in a given mode (default modes are "create", "edit" and "view") ; layouts in the "any" mode will be merged with layouts defined for specific modes.

The layout names refer to layouts defined on another extension point. Please see the layouts section for more information.

5.4.3.2. Deprecated configuration

Here is the old layout configuration that has been replaced by the above. If present, it is used instead of the new configuration for compatibility purposes.

```
<layout>
  <widget jsfcomponent="h:inputText"
    schemaname="dublincore" fieldname="title"
    required="true" />
  <widget jsfcomponent="h:inputTextarea"
    schemaname="dublincore" fieldname="description" />
```

```
  <widget jsfcomponent="h:inputText"
    schemaname="sample" fieldname="sample1" />
  <widget jsfcomponent="h:inputText"
    schemaname="sample" fieldname="sample2" />
</layout>
```

This section defines a series of widgets, that describe what the standard layout of this document type will be. A layout is a series of widgets, which make the association between the field of a schema with a JSF component.

The layout is used by the standard Nuxeo modification and summary views, to automatically display the document according to the layout rules. Note that the layout system is still young and is subject to minor changes in the future. Here we define four widgets, displayed as simple input fields or as a text area.

5.4.4. Containment rules

```
<type id="Folder" coretype="Folder">
  <subtypes>
    <type>Sample</type>
  </subtypes>
</type>
<type id="Workspace" coretype="Workspace">
  <subtypes>
    <type>Sample</type>
  </subtypes>
</type>
```

This contributes a rule to an already existing type, `Folder`. The subtypes element specifies which types can be created inside the type in which the element is embedded.

Here, we specify that our `Sample` type can be created in a `Folder` and a `Workspace`.

5.4.5. Summary

The final `OSGI-INF/ecm-types-contrib.xml` looks like:

```
<?xml version="1.0"?>
<component name="org.nuxeo.project.sample.ecm.types">
  <extension target="org.nuxeo.ecm.platform.types.TypeService" point="types">
    <type id="Sample" coretype="Sample">
      <label>Sample document</label>
      <icon>/icons/file.gif</icon>
      <default-view>view_documents</default-view>
      <layout>
        <widget jsfcomponent="h:inputText"
          schemaname="dublincore" fieldname="title"
          required="true" />
        <widget jsfcomponent="h:inputTextarea"
          schemaname="dublincore" fieldname="description" />
        <widget jsfcomponent="h:inputText"
          schemaname="sample" fieldname="sample1" />
        <widget jsfcomponent="h:inputText"
          schemaname="sample" fieldname="sample2" />
      </layout>
    </type>
    <type id="Folder" coretype="Folder">
      <subtypes>
        <type>Sample</type>
      </subtypes>
    </type>
    <type id="Workspace" coretype="Workspace">
      <subtypes>
        <type>Sample</type>
      </subtypes>
    </type>
  </extension>
</component>
```

Part II. Platform Services

Chapter 6. Exception Handling

6.1. Introduction

The ExceptionHandlingService handles the exceptions that bubbled outside of Nuxeo. It allows to define handlers to forward to an error page with adequate messages and logs.

It is composed of 3 extension points: the errorHandlers configures the handlers, the requestdump configures the way the request will be written to error log and listener allows to set a listener to the process so state can be restore in case of error.

This service is only available in Nuxeo 5.2

6.2. Extension Points

6.2.1. requestdump

The request dump extension point allows to contribute a RequestDumper class. It takes a request as parameter and output the string that will be written to error log. The default contribution is:

```
<extension
  target="org.nuxeo.ecm.platform.web.common.exceptionhandling.service.ExceptionHandlingService"
  point="requestdump">
  <requestdump
    class="org.nuxeo.ecm.platform.web.common.exceptionhandling.service.DefaultRequestDumper" />
</extension>
```

It writes the attributes of the request.

6.2.2. listener

The listener extension point allows to contribute a listener that will be called during the different phase of the exception handling. The default contribution does nothing (this is use by WebEngine):

```
  <extension
    target="org.nuxeo.ecm.platform.web.common.exceptionhandling.service.ExceptionHandlingService"
    point="listener">
    <listener
      class="org.nuxeo.ecm.platform.web.common.exceptionhandling.service.NullExceptionHandlingListener" />
  </extension>
```

This contribution is over written in the module nuxeo-platform-webapp-core with:

```
<extension
  target="org.nuxeo.ecm.platform.web.common.exceptionhandling.service.ExceptionHandlingService"
  point="listener">
  <listener
    class="org.nuxeo.ecm.webapp.shield.SeamExceptionHandlingListener" />
</extension>
```

This listener restore the faces/seam context before to write the attributes (use by NuxeoEP).

6.2.3. errorhandlers

The errorhandlers extension point allows to configure the different handler. General attribute are:

errorhandlers attributes

bundle
The message bundle to use. The default bundle used in Nuxeo is messages.

loggerName
The name of the logger to use. Default is nuxeo-error-log. If you change this value, you need to update your logging configuration.

defaultpage
The page to which you will be forwarded if no page is defined in the handler configuration.

The attributes for each handler are:

error
A regular expression, if matched the Exception we are handling, this handler will be used. The last handler should use ".*" as a Regular Expression.

message
The key, in the message bundle defined in the general attribute, of the message that will be output to the user.

page
The error page we should forward to. If not defined we will forward to the default error page as defined in the general attribute.

code
The status code set to the response. If none the default code will be set.

The default configuration is:

```
<extension
    target="org.nuxeo.ecm.platform.web.common.exceptionhandling.service.ExceptionHandlingService"
    point="errorhandlers">
   <errorHandlers bundle="messages" loggerName="nuxeo-error-log" defaultpage="nuxeo_error.jsp">
      <handlers>
        <handler error=".*NoSuchDocumentException" message="Error.Document.Not.Found" code="404"/>
        <handler error="javax.jcr.ItemNotFoundException" message="Error.Document.NotFound" code="404"/>
        <handler error=".*NoSuchPropertyException" message="Error.Document.NoSuchProperty"/>
        <handler error=".*SecurityException" message="Error.Insuffisant.Rights"/>
        <handler error=".*" message="Error.Unknown"/>
      </handlers>
    </errorHandlers>
  </extension>
```

Chapter 7. Actions, Views, Navigation URLs and JSF tags

7.1. Introduction

The User Interface framework uses different kinds of concepts to make the interface configurable in the same way that the application itself is.

Links and pages that appear on the site can be the result of a "static" template written in the XHTML language, but can also be the result of a configuration change, taken into account thanks to more "generic" XHTML pages.

The following chapters will explain how configuration and templating combine to achieve the UI of the site.

7.2. Actions

7.2.1. Concepts

In this chapter, an action will stand for any kind of command that can be triggered via user interface interaction. In other words, it will describe a link and other information that may be used to manage its display (the link label, an icon, security information for instance).

7.2.2. Manage actions

Custom actions can be contributed to the actions service, using its extension point. Their description is then available through this service to control where and how they will be displayed.

7.2.2.1. Register a new action

An action can be registered using the following example extension:

Example 7.1. Example of an action registration

```
<extension target="org.nuxeo.ecm.platform.actions.ActionService"
    point="actions">

  <action id="logout" link="#{loginLogoutAction.logout}"
    label="command.logout">
    <category>USER_SERVICES</category>
  </action>

</extension>
```

The above action will be used to display a "logout" link on the site. Here are its properties:

- `id`: string identifying the action. In the example, the action id is "logout".
- `label`: the action name that will be used when displaying the link. In the example, the label is "command.logout". This label is a message that will be translated at display.
- `link`: string representing the command the action will trigger. This string may represent a different action given the template that will display the action. In the example, a JSF command link will be used, so it represents an action method expression. The seam component called "loginLogoutAction" holds a method named "logout" that will perform the logout and return a string for navigation.
- `category`: a string useful to group actions that will be rendered in the same area of a page. An action can

define several categories. Here, the only category defined is "USER_SERVICES". It is designed to group all the actions that will be displayed on the right top corner of any page of the site.

Other properties can be used to define an action. They are listed here but you can have a look at the main actions contributions file for more examples: `nuxeo-platform-webapp-core/srs/main/resources/OSGI-INF/actions-contrib.xml`.

- `filter-ids`: id of a filter that will be used to control the action visibility. An action can have several filters: it is visible if all its filters grant the access.
- `filter`: a filter definition can be done directly within the action definition. It is a filter like others and can be referred by other actions.
- `icon`: the optional icon path for this action.
- `confirm`: an optional Javascript confirmation string that can be triggered when executing the command.
- `order`: an optional integer used to sort actions within the same category. This attribute may be deprecated in the future.
- `enabled`: boolean indicating whether the action is currently active. This can be used to hide existing actions when customizing the site behavior.

Actions extension point provides merging features: you can change an existing action definition in your custom extension point provided you use the same identifier. Properties holding single values (label, link for instance) will be replaced using the new value. Properties holding multiple values (categories, filters) will be merged with existing values.

7.2.2.2. Manage category to display an action at the right place

Actions in the same category are supposed to be displayed in the same area of a page. Here are listed the main default categories if you want to add an action there:

- `USER_SERVICES`: used to display actions in the top right corner of every page. The link attribute should look like a JSF action command link. See the example above.
- `VIEW_ACTION_LIST`: used for tabs displayed on every document. As each tab is not displayed in a different page, but just includes a specific template content in the middle of the page, the link attribute has to be a template path. For instance:

```
<action id="TAB_VIEW" link="/incl/tabs/document_view.xhtml" enabled="true"
    order="0" label="action.view.summary">
  <category>VIEW_ACTION_LIST</category>
  <filter-id>view</filter-id>
</action>

<action id="TAB_CONTENT" link="/incl/tabs/document_content.xhtml" order="10"
    enabled="true" label="action.view.content">
  <category>VIEW_ACTION_LIST</category>
  <filter-id>view_content</filter-id>
</action>
```

- `SUBVIEW_UPPER_LIST`: used to display actions just below a document tabs listing. As `USER_SERVICES`, these actions will be displayed using a command link, so the link attribute has to be an action method expression. For instance:

```
<action id="newSection" link="#{documentActions.createDocument('Section')}"
    enabled="true" label="command.create.section"
    icon="/icons/action_add.gif">
  <category>SUBVIEW_UPPER_LIST</category>
  <filter id="newSection">
    <rule grant="true">
      <permission>AddChildren</permission>
      <type>SectionRoot</type>
    </rule>
  </filter>
</action>

<action id="newDocument" link="select_document_type" enabled="true"
    label="action.new.document" icon="/icons/action_add.gif">
```

```
  <category>SUBVIEW_UPPER_LIST</category>
  <filter-id>create</filter-id>
</action>
```

7.2.2.3. Manage filters to control an action visibility

An action visibility can be controlled using filters. An action filter is a set of rules that will apply - or not - given an action and a context.

Filters can be registered using their own extension point, or registered implicitly when defining them inside of an action definition:

Example 7.2. Example of a filter registration

```
<filter id="view_content">
  <rule grant="true">
    <permission>ReadChildren</permission>
    <facet>Folderish</facet>
  </rule>
  <rule grant="false">
    <type>Root</type>
  </rule>
</filter>
```

Example 7.3. Example of a filter registration inside an action registration

```
<action id="newSection" link="#{documentActions.createDocument('Section')}"
    enabled="true" label="command.create.section"
    icon="/icons/action_add.gif">
  <category>SUBVIEW_UPPER_LIST</category>
  <filter id="newSection">
    <rule grant="true">
      <permission>AddChildren</permission>
      <type>SectionRoot</type>
    </rule>
  </filter>
</action>
```

A filter can accept any number of rules. It will grant access to an action if, among its rules, no denying rule (`grant=false`) is found and at least one granting rule (`grant=true`) is found. A general rule to remember is that if you would like to add a filter to an action that already has one or more filters, it has to hold constraining rules: a granting filter will be ignored if another filter is already too constraining.

If no rule is set, the filter will grant access by default.

The default filter implementation uses filter rules with the following properties:

- `grant`: boolean indicating whether this is a granting rule or a denying rule.
- `permission`: permission like "Write" that will be checked on the context for the given user. A rule can hold several permissions: it applies if user holds at least one of them.
- `facet`: facet like "Folderish" that can be set on the document type (`org.nuxeo.ecm.core.schema.types.Type`) to describe the document type general behavior. A rule can hold several facets: it applies if current document in context has at least one of them.
- `condition`: EL expression that can be evaluated against the context. The Seam context is made available for conditions evaluation. A rule can hold several conditions: it applies if at least one of the conditions is verified.
- `type`: document type to check against current document in context. A rule can hold several types: it

applies if current document is one of them. The fake 'Server' type is used to check the server context.

- `schema`: document schema to check against current document in context. A rule can hold several schemas: it applies if current document has one of them.

Filters do not support merging, so if you define a filter with an id that is already used in another contribution, only the first contribution will be taken into account.

7.2.2.4. Adapt templates to display an action

It is important to understand that an action does *not* define the way it will be rendered: This is left to pages, templates and other components displaying it. Most of the time, actions will be rendered as command links or command buttons.

For instance, actions using the `USER_SERVICES` category will be rendered as action links:

```
<nxu:methodResult name="actions"
    value="#{webActions.getActionsList('USER_SERVICES')}">
  <nxu:dataList layout="simple" var="action" value="#{actions}"
    rowIndexVar="row" rowCountVar="rowCount">
    <h:outputText value=" | " rendered="#{row!=(rowCount-1)}" />
    <nxh:commandLink action="#{action.getLink()}">
      <t:htmlTag value="br" rendered="#{row==(rowCount-1)}" />
      <h:outputText value="#{messages[action.label]}" />
    </nxh:commandLink>
  </nxu:dataList>
</nxu:methodResult>
```

The `nxu:methodResult` tag is only used to retrieve the list of actions declared for the `USER_SERVICES` category. The `nxh:commandLink` is used instead of a simple `h:commandLink` so that it executes commands that where described as action expression methods.

Another use case is the document tabs: actions using the `VIEW_ACTION_LIST` category will be rendered as action links too, but actions are managed by a specific seam component that will hold the information about the selected tab. When clicking on an action, this selected tab will be changed and the link it points to will be displayed.

7.3. Views

7.3.1. UI Views

First of all, we have to make the difference between a view in a standard JSF way (navigation case view id, navigation case output) and views in Nuxeo 5 (document type view, creation view)

7.3.2. Manage views

7.3.2.1. Standard JSF navigation concepts

A standard JSF navigation rule can be defined in the `OSGI-INF/deployment-fragment.xml` files, inside the `faces-config#NAVIGATION` directive.

Example 7.4. Example of a navigation rule case definitions

```
<extension target="faces-config#NAVIGATION">

    <navigation-case>
      <from-outcome>create_document</from-outcome>
      <to-view-id>/create_document.xhtml</to-view-id>
      <redirect />
    </navigation-case>

    <navigation-case>
```

```
      <from-outcome>view_documents</from-outcome>
      <to-view-id>/view_documents.xhtml</to-view-id>
      <redirect />
    </navigation-case>

</extension>
```

7.3.2.2. Nuxeo 5 views

A certain Nuxeo document type, can have defined a default view (used to view/edit the document) and a create view (used to create the document). These views are specified in the `OSGI-INF/ecm-types-contrib.xml` file, as in the following example.

Example 7.5. Example of view definitions for a document type

```
<extension target="org.nuxeo.ecm.platform.types.TypeService" point="types">
    <type id="Workspace" coretype="Workspace">
      <label>Workspace</label>
      <icon>/icons/workspace.gif</icon>
      <icon-expanded>/icons/workspace_open.gif</icon-expanded>
      <default-view>view_documents</default-view>
      <create-view>create_workspace</create-view>
    </type>
</extension>
```

The default view of a document is rendered as a list of tabs. As mentioned before, the document tabs are defined as actions in the `OSGI-INF/actions-contrib` file, having as category `VIEW_ACTION_LIST`. A tab can be added to a document default view as shown in the following example.

Example 7.6. Example of tab definition for a default view of a document type

```
<extension target="org.nuxeo.ecm.platform.actions.ActionService" point="actions">
    <action id="TAB_EDIT" link="/incl/tabs/document_edit.xhtml" enabled="true"
      order="20" label="action.view.edit" icon="/icons/file.gif">
      <category>VIEW_ACTION_LIST</category>
      <filter-id>edit</filter-id>
      <filter-id>mutable_document</filter-id>
    </action>
</extension>
```

7.4. Navigation URLs

There are two services that help building GET URLs to restore a Nuxeo context. The default configuration handle restoring the current document, the view, current tab and current sub tab.

7.4.1. Document view codec service

The service handling document views allows registration of codecs. Codecs manage coding of a document view (holding a document reference, repository name as well as key-named string parameters) in to a URL, and decoding of this URL into a document view.

Example 7.7. Example of a document view codec registration

```
<extension
  target="org.nuxeo.ecm.platform.url.service.DocumentViewCodecService"
  point="codecs">

  <documentViewCodec name="docid" enabled="true" default="true" prefix="nxdoc"
```

```
    class="org.nuxeo.ecm.platform.url.codec.DocumentIdCodec" />
  <documentViewCodec name="docpath" enabled="true" default="false" prefix="nxpath"
    class="org.nuxeo.ecm.platform.url.codec.DocumentPathCodec" />

</extension>
```

In this example, the docid codec uses the document uid to resolve the context. Urls are of the form http://site/nuxeo/nxdoc/demo/docuid/view. The docpath codec uses the document path to resolve the context. Urls are of the form http://site/nuxeo/nxpath/demo/path/to/my/doc@view.

Additional parameters are coded/decoded as usual request parameters.

Note that when building a document view, the url service will require a view id. The other information (document location and parameters) are optional, as long as they're not required for your context to be initialized correctly.

7.4.2. URL policy service

The service handling URLs allows registration of patterns. These patterns help saving the document context and restoring it thanks to information provided by codecs. The URL service will iterate through its patterns, and use the first one that returns an answer (proving decoding was possible).

Example 7.8. Example of a url pattern registration

```
<extension target="org.nuxeo.ecm.platform.ui.web.rest.URLService"
  point="urlpatterns">

  <urlPattern name="default" enabled="true">
    <defaultURLPolicy>true</defaultURLPolicy>
    <needBaseURL>true</needBaseURL>
    <needRedirectFilter>true</needRedirectFilter>
    <needFilterPreprocessing>true</needFilterPreprocessing>
    <codecName>docid</codecName>
    <actionBinding>#{restHelper.initContextFromRestRequest}</actionBinding>
    <documentViewBinding>#{restHelper.documentView}</documentViewBinding>
    <newDocumentViewBinding>#{restHelper.newDocumentView}</newDocumentViewBinding>
    <bindings>
      <binding name="tabId">#{webActions.currentTabId}</binding>
      <binding name="subTabId">#{webActions.currentSubTabId}</binding>
    </bindings>
  </urlPattern>

</extension>
```

In this example, the "default" pattern uses the above "docid" codec. Its is set as the default URL policy, so that it's used by default when caller does not specify a pattern to use. It needs the base URL: the docid codec only handles the second part if the URL. It needs redirect filter: it will be used to provide the context information to store. It needs filter preprocessing: it will be used to provide the context information to restore. It's using the docid codec.

The action binding method handles restoring of the context in the Seam context. It takes a document view as parameter. It requires special attention: if you're using conversations (as Nuxeo does by default), you need to annotate this method with a "@Begin" tag so that it uses the conversation identifier passed as a parameter if it's still valid, or initiates a new conversation in other cases. The method also needs to make sure it initializes all the seam components it needs (documentManager for instance) if they have not be in intialized yet.

The additional document view bindings are used to pass document view information through requests. The document view binding maps to corresponding getters and setters. The new document view binding is used to redirect to build GET URL in case request is a POST: it won't have the information in the URL so it needs to rebuild it.

Other bindings handle additional request parameters. In this example, they're used to store and restore tab and sub tab information (getters and setters have to be defined accordingly).

7.4.3. Additional configuration

The URL patterns used need to be registered on the authentication service so that they're considered as valid urls. Valid urls will be stored in the request, so that if authentication is required, user is redirected to the url after login.

Example 7.9. Example of a start url pattern registration

```
<extension
  target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
  point="startURL">

  <startURLPattern>
    <patterns>
      <pattern>nxdoc/</pattern>
    </patterns>
  </startURLPattern>

</extension>
```

Just the start of the url is required in this configuration. Contributions are merged: it is not possible to remove an existing start pattern.

The URL patterns used also need to be handled by the default nuxeo authentication service so that login mechanism (even for anonymous) applies for them.

Example 7.10. Example authentication filter configuration

```
<extension target="web#STD-AUTH-FILTER">
  <filter-mapping>
    <filter-name>NuxeoAuthenticationFilter</filter-name>
    <url-pattern>/nxdoc/*</url-pattern>
    <dispatcher>REQUEST</dispatcher>
    <dispatcher>FORWARD</dispatcher>
  </filter-mapping>
</extension>
```

This is a standard filter mapping configuration.

7.4.4. URL JSF tags

There are some JSF tags and functions helping you to define what kind of GET URL should be displayed on the interface.

Example 7.11. Example of nxd:restDocumentLink use

```
<nxd:restDocumentLink document="#{doc}">
  <nxh:outputText value="#{nxd:titleOrId(doc)}" />
</nxd:restDocumentLink>
```

In this example, the tag will print a simple link, using the default pattern, and build the document view using given document model, using its default view.

Please refer to the tag library documentation available at http://doc.nuxeo.org/5.1/tlddoc/nxd/restDocumentLink.html for additional parameters: it's possible to set the tab, sub tab, and use a specific URL pattern.

Note that you can also use JSF functions to build the GET URL. This is what's done for file links: the function queries the URL policy service to build the URL.

Example 7.12. Example of a jsf function use

```
<nxh:outputLink rendered="#{doc.hasSchema('file') and !empty doc.file.content}"
  value="#{nxd:fileUrl('downloadFile', doc, 'file:content', doc.file.filename)}">
  <nxh:graphicImage value="/icons/download.png" style="vertical-align:middle"
    title="#{doc.file.filename}" />
</nxh:outputLink>
```

Example 7.13. fileUrl method code

```
public static String fileUrl(String patternName, DocumentModel doc,
        String blobPropertyName, String filename) {
    try {
        DocumentLocation docLoc = new DocumentLocationImpl(doc);
        Map<String, String> params = new HashMap<String, String>();
        params.put(DocumentFileCodec.FILE_PROPERTY_PATH_KEY,
                blobPropertyName);
        params.put(DocumentFileCodec.FILENAME_KEY, filename);
        DocumentView docView = new DocumentViewImpl(docLoc, null, params);

        // generate url
        URLPolicyService service = Framework.getService(URLPolicyService.class);
        if (patternName == null) {
            patternName = service.getDefaultPatternName();
        }
        return service.getUrlFromDocumentView(patternName, docView,
                BaseURL.getBaseURL());

    } catch (Exception e) {
        log.error("Could not generate url for document file", e);
    }

    return null;
}
```

Similar methods exist for more complex urls, when handling files in list for instance. Please refer to the list at http://doc.nuxeo.org/5.1/tlddoc/nxd/tld-summary.html.

7.5. Nuxeo Document Lists Manager

The Document List Manager provides a service to manage lists of Nuxeo documents.

These lists of documents can have properties such as:

- a name, defined by `name` attribute.
- a scope (session or conversation), defined by `<isSession/>` tag - it defines if the memory storage occurs in the Seam session context or in the Seam conversation context.
- a persistence (SQL directory or not present), defined by `<persistent/>` tag - the service persists only the list of the document references, not the real documents; the lists of document references is persisted in a SQL directory, which is generic and does not need any configuration.

The lists of documents can be invalidated when Seam events are raised. This is usefull, for example, for resetting `CURRENT_SELECTION` lists when the user change the current folder or when a new search is performed.

Documents lists can be defined like in the following example (`OSGI-INF/documentslists-contrib.xml`):

Example 7.14. Example of documents lists definition

```
<extension target="org.nuxeo.ecm.webapp.documentsLists.DocumentsListsService" point="list">
```

```
    <documentsList name="CLIPBOARD">
      <category>CLIPBOARD</category>
      <imageURL>/img/clipboard.gif</imageURL>
      <title>workingList.clipboard</title>
      <defaultInCategory>false</defaultInCategory>
      <supportAppends>false</supportAppends>
    </documentsList>

    <documentsList name="CURRENT_SELECTION">
      <events>
        <event>folderishDocumentSelectionChanged</event>
        <event>searchPerformed</event>
      </events>
      <isSession>false</isSession>
    </documentsList>
</extension>
```

7.6. Nuxeo File Manager

The File Manager provides a service which exposes a simple API in order to create a Nuxeo document model from a passed blob file. Also, this service exposes as extension point which can be used to define plugins.

When the File Manager service is called, it will detect the mime-type of the passed blob, and will try to find a plugin to hadle the creation of a Nuxeo document model based on the detected mime-type.

Typical usages:

- txt/html/xml files - a `Note` document is created.
- image files - a `Picture` document is created (if `nuxeo-platform-imaging` addon is deployed).
- folder - a `Folder` document is created.
- other files - a `File` document is created.

Pugins can be defined like in the following example (`OSGI-INF/nxfilemanager-plugins-contrib.xml`):

Example 7.15. Example of File Manager plugins definition

```
<extension target="org.nuxeo.ecm.platform.filemanager.service.FileManagerService" point="plugins">
    <plugin name="Noteplugin"
      class="org.nuxeo.ecm.platform.filemanager.service.extension.NotePlugin">
      <filter>text/plain</filter>
      <filter>text/html</filter>
      <filter>application/xhtml+xml</filter>
      <filter>application/xml</filter>
      <filter>text/xml</filter>
    </plugin>
</extension>
```

As a client of the File Manager service can be used the browser plugins (for Firefox and Internet Explorer) which can be be downloaded through the links from the default Nuxeo 5 login page. These plugins enable the user to create Nuxeo documents just by dragging & dropping folders/files to the browser. The plugins use a restlet (HTTP API) to send files/folders to the Nuxeo 5 Platform. The restlets use the File Manager serivce in order to create a Nuxeo document from the passed file.

7.7. Nuxeo JSF tags

Please refer to the tag library documentation available at http://doc.nuxeo.org/5.1/tlddoc/.

Chapter 8. Layouts

Let our artists go wild on imaginative page layouts.

-- Grant Morrison

8.1. Introduction

Layouts are used to generate pages rendering from an xml configuration.

In a document oriented perspective, layouts are mostly used to display a document metadata in different use cases: present a form to set its schemas fields when creating or editing the document, and present these fields values when simply displaying the document. A single layout definition can be used to address these use cases as it will be rendered for a given document and in a given mode.

In this chapter we will see how to define a layout, link it to a document type, and use it in XHTML pages.

8.1.1. Layouts

A *layout* is a group of widgets that specifies how widgets are assembled and displayed. It manages widget rows and has global control on the rendering of each of its widgets.

8.1.2. Widgets

It's all the same machine, right? The Pentagon, multinational corporations, the police! You do one little job, you build a widget in Saskatoon and the next thing you know it's two miles under the desert, the essential component of a death machine!

-- Holloway, Cube

A *widget* defines how one or several fields from a schema will be presented on a page. It can be displayed in several modes and holds additional information like for instance the field label. When it takes user entries, it can perform conversion and validation like usual JSF components.

8.1.3. Widget types

A widget definition includes the mention of its *type*. Widget types make the association between a widget definition and the jsf component tree that will be used to render it in a given mode.

8.1.4. Modes

The *layout modes* can be anything although some default modes are included in the application: *create*, *edit*, *view* and *search*.

The widget modes are more restricted and widget types will usually only handle two modes: edit and view. The widget mode is computed from the layout mode following this rule: if the layout is in mode create, edit or search, the widget will be in edit mode. Otherwise the widget will be in view mode.

It is possible to override this behaviour in the widget definition, and state that, for instance, whatever the layout mode, the widget will be in view mode so that it only displays read-only values. The pseudo-mode "hidden" can also be used in a widget definition to exclude this widget from the layout in a given mode.

The pseudo mode *"any*" is only used in layouts and widgets definitions to set up default values.

8.2. Manage layouts

Custom layouts can be contributed to the web layout service, using its extension point. The layout definition is then available through the service to control how it will be displayed in a given mode.

Some jsf tags have been added to the Nuxeo ECM layout tag library to make then easily available from an xhtml page.

8.2.1. Layout registration

Layouts are registered using a regular extension point on the Nuxeo ECM layout service. Here is a sample contribution.

Example 8.1. Sample layout contribution to the layout service.

```
<?xml version="1.0"?>

<component name="org.nuxeo.ecm.platform.forms.layouts.webapp">

  <extension target="org.nuxeo.ecm.platform.forms.layout.WebLayoutManager"
    point="layouts">

    <layout name="heading">
      <templates>
        <template mode="any">/layouts/layout_default_template.xhtml</template>
      </templates>
      <rows>
        <row>
          <widget>title</widget>
        </row>
        <row>
          <widget>description</widget>
        </row>
      </rows>
      <widget name="title" type="text">
        <labels>
          <label mode="any">label.dublincore.title</label>
        </labels>
        <translated>true</translated>
        <fields>
          <field>dc:title</field>
        </fields>
        <properties widgetMode="edit">
          <property name="required">true</property>
        </properties>
      </widget>
      <widget name="description" type="textarea">
        <labels>
          <label mode="any">label.dublincore.description</label>
        </labels>
        <translated>true</translated>
        <fields>
          <field>dc:description</field>
        </fields>
      </widget>
    </layout>

  </extension>

</component>
```

8.2.2. Layout definition

The above layout definition is used to display the title and the description of a document. Here are its properties:

- `name`: string used as an identifier. In the example, the layout name is "heading".
- `templates`: list of templates to use for this layout global rendering. In the example, the layout template in any mode is the xhtml file at "/layouts/layout_default_template.xhtml". Please refer to section about custom layout templates for more information.

- `rows`: definition about what widgets will have to be displayed on this row. Each row can hold several widgets, and an empty widget tag can be used to control the alignment. The widget has to match a widget name given in this layout definition. In the example, two rows have been defined, the first one will hold the "title" widget, and the second one will hold the "description" widget.

- `widget`: a layout definition can hold any number of widget definitions. If the widget is not referenced in the rows definition, it will be ignored. Since 5.1.7 and 5.2.0, it will be searched in the global widget registry before being ignored. This new feature is a convenient way to share widget definitions between layouts. Please refer the widget definition section.

8.2.3. Widget definition

Two widget definitions are presented on the above example. Let's look into the "title" widget and present its properties:

- `name`: string used as an identifier in the layout context. In the example, the widget name is "title".

- `type`: the widget type that will manage the rendering of this widget. In this example, the widget type is "text". This widget type is a standard widget types, more information about widget types is available here.

- `labels`: list of labels to use for this widget in a given mode. If no label is defined in a specific mode, the label defined in the "any" mode will be taken as default. In the example, a single label is defined for any mode to the "label.dublicore.title" message. If no label is defined at all, a default label will be used following the convention: "label.widget.[layoutName].[widgetName]".

- `translated`: string representing a boolean value ("true" or "false") and defaulting to "false". When set as translated, the widget labels will be treated as messages and displayed translated. In the example, the "label.dublincore.title" message will be translated at rendering time. Default is true.

- `fields`: list of fields that will be managed by this widget. In the example, we handle the field "dc:title" where "dc" is the prefix for the "dublincore" schema. If the schema you would like to use does not have a prefix, use the schema name instead. Note that most of standard widget types only handle one field. Side note: when dealing with an attribute from the document that is not a metadata, you can use the property name as it will be resolved like a value expression of the form #{document.attribute}.

- `properties`: list of properties that will apply to the widget in a given mode. Properties listed in the "any" mode will be merged with properties for the specific mode. Depending on the widget type, these properties can be used to control what jsf component will be used and/or what attributes will be set on these components. In standard widget types, only one component is used given the mode, and properties will be set as attributes on the component. For instance, when using the "text" widget type, every property accepted by the "<h:inputText />" tag can be set as properties on "edit" and "create" modes, and every property accepted by the "<h:outputText />" tag can be set as properties. Properties can also be added in a given widget mode.

Additional properties can be set on a widget:

- `helpLabels`: list that follows the same pattern as labels, but used to set help labels.

- `widgetModes`: list of local modes used to override the local mode (from the layout).

- `subWidgets`: list of widget definitions, as the widget list, used to describe sub widgets use to help the configuration of some complex widget types.

Here is a more complex layout contribution that shows the syntax to use for these additional properties:

Example 8.2. Sample complex layout contribution to the layout service.

```
<?xml version="1.0"?>

<component name="org.nuxeo.ecm.platform.forms.layouts.webapp">

  <!-- WARNING: this extension point is only available in 5.1.7 and 5.2.0 -->
  <extension target="org.nuxeo.ecm.platform.forms.layout.WebLayoutManager"
    point="widgets">

    <!-- global definition of a widget so that it can be used
      in several layouts -->
    <widget name="description" type="textarea">
      <labels>
        <label mode="any">description</label>
      </labels>
      <translated>true</translated>
      <fields>
        <field>dc:description</field>
      </fields>
      <properties widgetMode="edit">
        <property name="styleClass">dataInputText</property>
      </properties>
    </widget>

  </extension>

  <extension target="org.nuxeo.ecm.platform.forms.layout.WebLayoutManager"
    point="layouts">

    <layout name="complex">
      <templates>
        <template mode="any">/layouts/layout_default_template.xhtml</template>
      </templates>
      <rows>
        <row>
          <widget>identifier</widget>
        </row>
        <row>
          <!-- reference a global widget -->
          <widget>description</widget>
        </row>
      </rows>
      <widget name="identifier" type="text">
        <labels>
          <label mode="any">label.dublincore.title</label>
        </labels>
        <translated>true</translated>
        <fields>
          <field>uid:uid</field>
        </fields>
        <widgetModes>
          <!-- not shown in create mode -->
          <mode value="create">hidden</mode>
        </widgetModes>
        <properties widgetMode="edit">
          <!-- required in widget mode edit -->
          <property name="required">true</property>
        </properties>
        <properties mode="view">
          <!-- property applying in view mode -->
          <property name="styleClass">cssClass</property>
        </properties>
      </widget>
    </layout>

  </extension>

</component>
```

8.2.4. EL expressions in layouts and widgets

TODO: explain what variables are available in the context

8.3. Document layouts

Layouts can be linked to a document type definition by specifying the layout names:

```
<layouts mode="any">
  <layout>heading</layout>
  <layout>note</layout>
</layouts>
```

Layouts are defined in a given mode; layouts in the "any" mode will be used as default when no layouts are given in specific modes.

Since 5.2.GA, it is possible to merge layouts when redefining the document type, adding a property append="true":

```
<layouts mode="any" append="true">
  <layout>newLayout</layout>
</layouts>
```

8.4. Layout display

Layouts can be displayed thanks to a series a JSF tags that will query the web layout service to get the layout definition and build it for a given mode.

For instance, we can use the `documentLayout` tag to display the layouts of a document:

```
<div xmlns="http://www.w3.org/1999/xhtml"
    xmlns:nxl="http://nuxeo.org/nxforms/layout">
  <nxl:documentLayout mode="view" value="#{currentDocument}" />
</div>
```

We can also display a specific layout for a document, even if it is not specified in the document type definition:

```
<div xmlns="http://www.w3.org/1999/xhtml"
    xmlns:nxl="http://nuxeo.org/nxforms/layout">
  <nxl:layout name="heading" mode="view" value="#{currentDocument}" />
</div>
```

Please refer to the tag library documentation available at http://doc.nuxeo.org/5.1/tlddoc/.

8.5. Standard widget types

A series of widget types has been defined for the most generic uses cases.

Please refer to the tag library documentation available at http://doc.nuxeo.org/5.1/tlddoc/ for nuxeo jsf tags.

8.5.1. text

The text widget displays an input text in create or edit mode, with additional message tag for errors, and a regular text output in any other mode. Widgets using this type can provide properties accepted on a `<h:inputText />` tag in create or edit mode, and properties accepted on a `<h:outputText />` tag in other modes.

8.5.2. int

The int widget displays an input text in create or edit mode, with additional message tag for errors, and a regular text output in any other mode. It uses a number converter. Widgets using this type can provide properties accepted on a `<h:inputText />` tag in create or edit mode, and properties accepted on a `<h:outputText />` tag in other modes.

8.5.3. secret

The secret widget displays an input secret text in create or edit mode, with additional message tag for errors, and nothing in any other mode. Widgets using this type can provide properties accepted on a `<h:inputSecret />` tag in create or edit mode.

8.5.4. textarea

The textarea widget displays a textarea in create or edit mode, with additional message tag for errors, and a regular text output in any other mode. Widgets using this type can provide properties accepted on a `<h:inputTextarea />` tag in create or edit mode, and properties accepted on a `<h:outputText />` tag in other modes.

8.5.5. datetime

The datetime widget displays a javascript calendar in create or edit mode, with additional message tag for errors, and a regular text output in any other mode. It uses a date time converter. Widgets using this type can provide properties accepted on a `<nxu:inputDatetime />` tag in create or edit mode, and properties accepted on a `<h:outputText />` tag in other modes. The converter will also be given these properties.

8.5.6. template

The template widget displays a template content whatever the mode. Widgets using this type must provide the path to this template; this template can check the mode to adapt the rendering.

Information about how to write a template is given in the custom widget template section.

8.5.7. file

The file widget displays a file uploader/editor in create or edit mode, with additional message tag for errors, and a link to the file in other modes. Widgets using this type can provide properties accepted on a `<nxu:inputFile />` tag in create or edit mode, and properties accepted on a `<nxu:outputFile />` tag in other modes.

8.5.8. htmltext

The htmltext widget displays an html text editor in create or edit mode, with additional message tag for errors, and a regular text output in other modes (without escaping the text). Widgets using this type can provide properties accepted on a `<nxu:editor />` tag in create or edit mode, and properties accepted on a `<nxu:outputText />` tag in other modes.

8.5.9. selectOneDirectory

The selectOneDirectory widget displays a selection of directory entries in create or edit mode, with additional message tag for errors, and the directory entry label in other modes. Widgets using this type can provide properties accepted on a `<nxd:selectOneListbox />` tag in create or edit mode, and properties accepted on a `<nxd:directoryEntryOutput />` tag in other modes.

8.5.10. selectManyDirectory

The selectManyDirectory widget displays a multi selection of directory entries in create or edit mode, with additional message tag for errors, and the directory entries labels in other modes. Widgets using this type can provide properties accepted on a `<nxd:selectManyListbox />` tag in create or edit mode, and properties accepted on a `<nxd:directoryEntryOutput />` tag in other modes.

8.5.11. list

The list widget displays an editable list of items in create or edit mode, with additional message tag for errors, and the same list of items in other modes. Items are defined using sub wigdets configuration. This actually a template widget type whose template uses a `<nxu:inputList />` tag in edit or create mode, and a table iterating over items in other modes.

8.5.12. checkbox

The checkbox widget displays a checkbox in create, edit and any other mode, with additional message tag for errors. Widgets using this type can provide properties accepted on a `<h:selectBooleanCheckbox />` tag in create, edit mode, and other modes.

8.6. Custom templates

Some templating feature have been made available to make it easier to control the layouts and widgets rendering.

8.6.1. Custom layout template

A layout can define an xhtml template to be used in a given mode. Let's take a look at the default template structure.

Example 8.3. Default layout template

```
<f:subview
  xmlns:f="http://java.sun.com/jsf/core"
  xmlns:h="http://java.sun.com/jsf/html"
  xmlns:nxl="http://nuxeo.org/nxforms/layout"
  xmlns:nxu="http://nuxeo.org/nxweb/util"
  xmlns:nxd="http://nuxeo.org/nxweb/document">

  <f:subview rendered="#{layout.mode != 'edit' and layout.mode != 'create'}">

    <table class="dataInput">
      <tbody>

        <nxl:layoutRow>
          <tr>
            <nxl:layoutRowWidget>
              <td class="labelColumn">
                <h:outputText value="#{widget.label}" rendered="#{!widget.translated}" />
                <h:outputText value="#{messages[widget.label]}" rendered="#{widget.translated}" />
              </td>
              <td class="fieldColumn">
                <nxl:widget widget="#{widget}" mode="#{widget.mode}" value="#{value}" />
              </td>
            </nxl:layoutRowWidget>
          </tr>
        </nxl:layoutRow>

      </tbody>
    </table>

  </f:subview>

  <f:subview rendered="#{layout.mode == 'edit' or layout.mode == 'create'}">

    <table class="dataInput">
      <tbody>

        <nxl:layoutRow>
          <tr>
            <nxl:layoutRowWidget>
              <td class="labelColumn">
                <h:outputText value="#{widget.label}" rendered="#{!widget.translated}"
                  styleClass="#{nxu:test(widget.required, 'required', '')}" />
                <h:outputText value="#{messages[widget.label]}" rendered="#{widget.translated}"
                  styleClass="#{nxu:test(widget.required, 'required', '')}" />
              </td>
              <td class="fieldColumn">
                <nxl:widget widget="#{widget}" mode="#{widget.mode}" value="#{value}" />
              </td>
            </nxl:layoutRowWidget>
          </tr>
        </nxl:layoutRow>
```

```
      </tbody>
    </table>

  </f:subview>

</f:subview>
```

This template is intended to be unused in any mode, so the layout mode is checked to provide a different rendering in "edit" or "create" modes and other modes.

When this template is included in the page, several variables are made available:

- `layout`: the computed layout value ; its mode and number of columns can be checked on it.
- `value` or `document`: the document model (or whatever item used as value).

The layout system integration using facelets features requires that iterations are performed on the layout rows and widgets. The <nxl:layoutRow> and <nxl:layoutRowWidget /> trigger these iterations. Inside the layoutRow tag, two more variables are made available: `layoutRow` and `layoutRowIndex`. Inside the layoutRowWidget, two more variables are made available: `widget` and `widgetIndex`.

These variables can be used to control the layout rendering. For instance, the default template is the one applying the "required" style on widget labels, and translating these labels if the widget must be translated. It also makes sure widgets on the same rows are presented in the same table row.

8.6.2. Custom widget template

The template widget type makes it possible to set a template to use as an include.

Let's have a look at a sample template used to present contributors to a document.

Example 8.4. Sample template for a widget

```
<f:subview xmlns:f="http://java.sun.com/jsf/core"
  xmlns:h="http://java.sun.com/jsf/html"
  xmlns:a4j="https://ajax4jsf.dev.java.net/ajax"
  xmlns:t="http://myfaces.apache.org/tomahawk"
  xmlns:nxdir="http://nuxeo.org/nxdirectory">

  <t:dataList id="#{widget.id}" var="listItem" value="#{field_0}"
    layout="simple" styleClass="standardList">
    <h:graphicImage value="/icons/html.png" />
    <h:commandLink value="#{listItem}" immediate="true"
      action="#{userManagerActions.viewUser}">
      <f:param name="usernameParam" value="#{listItem}" />
    </h:commandLink>
    <br />
  </t:dataList>

</f:subview>
```

This widget presents the contributors of a document with specific links on each on these user identifier information.

Having a widget type just to perform this kind of rendering would be overkill, so using a widget with type "template" can be useful here.

When this template is included in the page, the `widget` variable is made available:

Some rules must be followed when writing xhtml to be included in templates:

- Use the widget id as identifier: the widget id is computed to be unique within the page, so it should be used instead of fixed id attributes so that another widget using the same template will not introduce duplicated ids in the jsf component tree.
- Use the variable with name following the `field_n` pattern to reference field values. For instance, binding a jsf component value attribute to `#{field_0}` means binding it to the first field definition.

8.6.3. Builtin templates to handle complex properties

8.6.3.1. List widget template

The standard widget type "list" is actually a widget of type "template" using a static template path: /widgets/list_widget_template.xhtml. If this default behaviour does not suit your needs, you can simply copy this template, make your changes, and use a widget of type "template" with the new template path.

This template assumes that each element of the list will be displayed using subwidgets definitions.

For instance, to handle a list of String elements, you can use the definition:

```
<widget name="contributors" type="list">
  <fields>
    <field>dc:contributors</field>
  </fields>
  <subWidgets>
    <widget name="contributor" type="text">
      <fields>
        <field></field>
      </fields>
    </widget>
  </subWidgets>
</widget>
```

The empty field definition in the subwidget is used to specify that each element of the list is itself the element to display.

To handle a list of complex properties (each entry of the list is a map with keys 'name' and 'email' for instance), you can use the definition:

```
<widget name="employees" type="list">
  <fields>
    <field>company:employees</field>
  </fields>
  <subWidgets>
    <widget name="name" type="text">
      <fields>
        <field>name</field>
      </fields>
    </widget>
    <widget name="email" type="text">
      <fields>
        <field>email</field>
      </fields>
    </widget>
  </subWidgets>
</widget>
```

8.6.3.2. Complex widget template

A builtin template has been added to handle complex properties. It is available at /widgets/complex_widget_template.xhtml. It assumes that each element of the complex property will be displated using subwidgets definitions.

To handle a complex property (the value is a map with keys 'name' and 'email' for instance, you can use the definition:

```
<widget name="manager" type="template">
  <fields>
    <field>company:manager</field>
  </fields>
  <properties mode="any">
    <property name="template">
      /widgets/complex_widget_template.xhtml
    </property>
  </properties>
  <subWidgets>
    <widget name="name" type="text">
      <fields>
        <field>name</field>
      </fields>
    </widget>
    <widget name="email" type="text">
      <fields>
        <field>email</field>
      </fields>
    </widget>
  </subWidgets>
</widget>
```

8.7. Custom widget types

Custom widget types can be added to the standard list thanks to another extension point on the web layout service.

Here is a sample widget type registration:

Example 8.5. Sample widget type contribution to the layout service.

```
<?xml version="1.0"?>

<component name="org.nuxeo.ecm.platform.forms.layout.MyContribution">

  <extension target="org.nuxeo.ecm.platform.forms.layout.WebLayoutManager"
    point="widgettypes">

    <widgetType name="customtype">
      <handler-class>
        org.myproject.MyCustomWidgetTypeHandler
      </handler-class>
      <property name="foo">bar</property>
    </widgetType>

  </extension>

</component>
```

The custom widget type class must follow the org.nuxeo.ecm.platform.forms.layout.facelets.WidgetTypeHandler interface.

Additional properties can be added to the type registration so that the same class can be reused with a different behaviour given the property value.

The widet type handler is used to generate facelet tag handlers dynamically taking into account the mode, and any other properties that can be found on a widget.

The best thing to do before writing a custom widget type handler is to go see how standard widget type handlers are implemented, as some helper methods can be reused to ease implementation of specific behaviours.

8.8. Generic layout usage

Layouts can be used with other kind of objects than documents.

The field definition has to match a document property for which setters and getters will be available, or the "value" property must be passed explicitely for the binding to happen. Depending on the widget, other kinds of bindings can be done.

Chapter 9. Event Listeners and Scheduling

9.1. Introduction

Events and event listeners have been introduced at the Nuxeo core level to allow pluggable behaviours when managing documents (or any kinds of objects of the site).

Whenever an event happens (document creation, document modification, relation creation, etc...), an event is sent to the event service that dispatches the notification to its listeners. Listeners can perform whatever action when receiving an event.

9.2. Concepts

A core event has a source which is usually the document model currently being manipulated. It can also store the event identifier, that gives information about the kind of event that is happening, as well as the principal connected when performing the operation, an attached comment, the event category, etc..

Events sent to the event service have to follow the `org.nuxeo.ecm.core.api.event.CoreEvent` interface.

A core event listener has a name, an order, and may have a set of event identifiers it is supposed to react to. Its definition also contains the operations it has to execute when receiving an interesting event.

Event listeners have to follow the `org.nuxeo.ecm.core.listener.EventListener` interface.

Several event listeners exist by default in the nuxeo platform, for instance:

- `DublincoreListener`: it listens to document creation/modification events and sets some dublincore metadata accordingly (date of creation, date of last modification, document contributors...)
- `DocUidGeneratorListener`: it listens to document creation events and adds an identifier to the document if an uid pattern has been defined for this document type.
- `DocVersioningListener`: it listens to document versioning change events and changes the document version numbers accordingly.

9.3. Adding an event listener

Event listeners can be plugged using extension points. Here are some examples of event listeners registration.

Example 9.1. DublincoreListener registration sample

```
<?xml version="1.0"?>
<component name="DublinCoreStorageService" version="1.0.0">
  <extension target="org.nuxeo.ecm.core.listener.CoreEventListenerService"
     point="listener">
    <listener name="dclistener"
      class="org.nuxeo.ecm.platform.dublincore.listener.DublinCoreListener"
      order="120" />
  </extension>
</component>
```

Example 9.2. UIDGenerator listener registration sample with event filtering

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.platform.uidgen.service.UIDGeneratorService">
  <extension target="org.nuxeo.ecm.core.listener.CoreEventListenerService"
      point="listener">
    <listener name="uidlistener"
        class="org.nuxeo.ecm.platform.uidgen.corelistener.DocUIDGeneratorListener"
        order="10">
      <eventId>documentCreated</eventId>
    </listener>
  </extension>
</component>
```

The only thing needed to add an event listener is to declare its name and its class. Sometimes the order in which listeners are called matters so an integer order can be set to control it. A filtering on event ids can be done when registering it too, though the notification method could handle it too.

For instance, the UIDgenerator service will only be notified when the event service receives a document creation event.

9.4. Upgrading an event listener

Since release of Nuxeo EP version 5.0 M3, events involving documents send the document model as source of the event. They used to send the document itself, which was wrong and has been changed in a compatible way.

Old school event listeners should still work ok for now, but should be migrated soon as the compatibility may introduce bugs and will be removed shortly.

To migrate your event listener, make it implement the empty interface `org.nuxeo.ecm.core.listener.DocumentModelEventListener`, and make it deal with a `DocumentModel` instead of a `Document` as event source.

If your event listener does not care about the source, or the event it deals with is not a document, you do not have to do anything.

9.5. Adding an event

To add an event, you have to create it and then notify listeners passing the even to the listener service. Here is a sample code on how to do it:

```
CoreEvent coreEvent = new CoreEventImpl(eventId, source, options,
       getPrincipal(), category, comment);

CoreEventListenerService service = NXCore.getCoreEventListenerService();

if (service != null) {
    service.notifyEventListeners(coreEvent);
} else {
    throw new ClientException("Can't get Event Listener Service");
}
```

9.6. From CoreEvents to JMS Messages

Events that are fired at the core level are forwarded to a JMS topic called NXPMessage.

This forwarding is done by a dedicated CoreEventListener (called JMSEventListener contributed by the `nuxeo-platform-events-core` bundle).

In order to be sure that when an JMS event is received the associated DocumentModel is available, all document oriented messages that may occur at core level are forwarded to the JMS topic when the session repository is saved (ie: when data is committed).

In some cases, depending on own the Core API is used, some messages can be duplicated within the same transaction (like modifying several times the same document), the JMSEventListener marks all duplicated messages before sending them to JMS, its JMS messages receiver to choose to process or not the duplicated messages.

During the forwarding on the JMS Topic, the coreEvents are converted to EventMessage. The main difference is that the EventMessage does not contains the DocumentData (ie: all schemas and fields are unloaded), this is done in order to avoid overloading JMS.

9.7. Adding a JMS message listener

The simplest way to add a JMS message listener is simply to define a Message Driven Bean that is bound to the NXPMessage Topic.

Here is a simple example a the definition of such a MDB :

```
@MessageDriven(activationConfig = {
  @ActivationConfigProperty(propertyName = "destinationType", propertyValue = "javax.jms.Topic"),
  @ActivationConfigProperty(propertyName = "destination", propertyValue = "topic/NXPMessages"),
  @ActivationConfigProperty(propertyName = "providerAdapterJNDI",
    propertyValue = "java:/NXCoreEventsProvider"),
  @ActivationConfigProperty(propertyName = "acknowledgeMode", propertyValue = "Auto-acknowledge") })
@TransactionManagement(TransactionManagementType.CONTAINER)
public class NXAuditMessageListener implements MessageListener {

    private static final Log log = LogFactory.getLog(NXAuditMessageListener.class);

    @TransactionAttribute(TransactionAttributeType.REQUIRED)
    public void onMessage(Message message) {
        try {

            final Serializable obj = ((ObjectMessage) message).getObject();
            if (!(obj instanceof DocumentMessage)) {
                log.debug("Not a DocumentMessage instance embedded ignoring.");
                return;
            }

            DocumentMessage docMessage = (DocumentMessage) obj;

            String eventId = docMessage.getEventId();
            log.debug("Received a message with eventId: " + eventId);
...
```

The `DocumentMessage` is a subclass of the `DocumentModel`.

An important point to remember is that the MDB is executed asynchronously in a dedicated thread:

- there is no JAAS Session established: you can not access the repository without this
- the `DocumentMessage` is not bound to an existing `CoreSession`: you can not use the `DocumenMessage` to do lazy loading (ie: `DocumentMessage.getProperty()`)

So, in order to extract some document oriented properties of the document associated to the event, you must:

- Establish a JAAS Session
- get a connected `DocumentModel` using the `DocumentRef` provided by the `DocumentMessage`

Here is a code sample for this:

```
LoginContext lc;
CoreSession session;
String repositoryName = docMessage.getRepositoryName();
try {
    log.debug("trying to connect to ECM platform");
    lc = Framework.login();
    session = Framework.getService(RepositoryManager.class).getRepository(repositoryName).open();
    DocumentModel connectedDoc = session.getDocument(docMessage.getRef());
...
} finally {
   if (session != null)
      CoreInstance.getInstance().close(session.getSessionId())
   if (lc != null)
```

```
        lc.logout();
}
```

9.8. Scheduling

XXX TODO: FG

Chapter 10. User Notification Service

10.1. Introduction

The notification framework provides a way to notify the users regarding different events as they happen in the system.

10.2. Notification concept

A notification is an alert that is sent to some users when an event takes place in the system (e.g. a document is created or deleted, a document is modified or published, a comment was entered, etc.).

A notification is defined by following attributes:

- name: must be unique
- channel: email, SMS, Jabber, etc.
- events: a series of events to which it responds
- template: is a file in which is stored the text that is sent to an user - it may be dynamic

In order to define a notification, one must declare all these attributes in a contribution file.

10.3. Notification channels

The channel is the communication channel used by the notification service to send alerts(notifications) to users.

In order to define a new channel (by default only email channel is used) the `ChannelNotificator` interface must be implemented.

It has 2 methods:

- `isInterestedInNotification(Notification docMessage)` usually checks if the channel is right.
- `sendNotification(DocumentMessage docMessage)` sends the actual notification. For now only email notification channel is implemented.

10.4. E-mail notifications

A notification must be defined in a xml file like the default `notification-contrib.xml` that is used by default.

Default defined notification may be disabled by setting *`enabled="false"`*.

In order to define a new notification one must place a definition like this one in his contribution file :

```
<notification name="Modification" channel="email" enabled="true"
              availableIn="workspace" autoSubscribed="false"
              template="modif" subject="Document modified"
              subjectTemplate="subjectModif"
              label="label.nuxeo.notifications.modif">
  <event name="documentModified"/>
  <event name="contentSubdocumentModified"/>
</notification>
```

As you may see above a notification must declare a list of events to which it reacts. In our case `documentModified` `contentSubdocumentModified`.

Also a notification must have a name that must be unique within the application. A label must also be specified for i18n.

The attribute enabled is used to enable / disable specific notifications.

The attribute *autoSubscribed* is set to true when we want that a notification is sent to all concerned users. In this case within the `eventInfo` map there must be loaded also the users that are concerned. For example if we want that some users (ex: administrators or workflow manager) to get a notification each time a task is assigned to them, we must use `autoSubscribed="true"` and put the usernames of all users in the `eventInfo` of the event under the key `recepients`.

The attribute *availableIn* is used in order to restrict the scope of a notification. For example if we want to define a notification that is triggered each time the document is modified, then it would not be used inside a section, because sections contain documents that cannot be modified, only published. So in order to hide this notification inside a section, we specify *availableIn="workspace"*. The accepted values are *workspace*, *section* and *all*.

The template attribute specifies the name of a template that will be used for generating the body of the email(notification). This name is associated with a file using another extension point like this: `<template name="modif" src="templates/modif.ftl" />`.

Inside a `*.ftl` file there may be inserted some dynamic parts like the document name, the user triggering the event, etc. Any data that one wishes to put inside the body of the email, or its subject, he must put that data in the `eventInfo` map under a unique key. Then inside the template file that data will be displayed using *${key}*.

For the email notification a subject is used. This subject is a string but is also dynamic following the same rules as the body inside the template files. For those who wants more parameterizable subject, you can use the subjectTemplate attribute : it specfiies the name of the template used for generating the subject and into which you can write dynamic content. As for the template attribute, the association between the name and the file is done using an extension point like `<template name="subjectModif" src="templates/subjectModificationTemplate.ftl" />`.

If both subject and subjectTemplate attributes are filled, the subjectTemplate attribute will be used to generate the subject.

Note that if you are writing an HTML based template it will be advised to use HTML encoded letters when there is accentuated letters (in french for example "é" will be "é"). The *htmlEscape* method is provided while writing templates to transform accentued characters in data from the `eventInfo` map.

```
<p> Your comment here : ${htmlEscape(key)} </p>
```

Chapter 11. Indexing & Searching

This chapter presents the architecture of the indexing and search service in Nuxeo EP 5.

11.1. Introduction

This chapter is under construction. XXX TODO: GR+JA

11.2. Configuration

For obvious performance and volume considerations, the search service doesn't index all the content of the application, nor does it provide the full content in search results. This must be specified in the configuration, along with what to do with the data.

The search service configuration is done by the standard Nuxeo Runtime extension point system. The target is `org.nuxeo.ecm.core.search.service` For a complete example, check the default configuration file: `nuxeo-platform-search-core/OSGI_INF/nxsearch-contrib.xml` (from Nuxeo EP sources).

11.2.1. Concepts

The main concepts are *Resource* and *Field*. A Resource holds several fields. It has a name and a prefix, which is used, e.g, in queries. Resources are supposed to be unsplittable, but they are usually aggregated. It's up to the backend to decide how to handle that aggregation; this goes beyond the scope of the present documentation.

At this point the search service handles documents only, so it's safe to say that documents correspond to aggregates of resources, and resources correspond to schemas. In the future, there'll be more types of resources.

11.2.2. The `indexableDocType` extension point

The core types of documents to index have to be registered against this extension point, in which the schemas to index are bound to each document type.

Here's an example demonstrating the available patterns:

```
<extension target="org.nuxeo.ecm.core.search.service.SearchServiceImpl"
           point="indexableDocType">
(...)
  <indexableDocType name="DocType1" indexAllSchemas="true"/>
  <indexableDocType name="DocType2" indexAllSchemas="true">
    <exclude>unwanted_schema</exclude>
  </indexableDocType>
  <indexableDocType name="DocType3">
    <resource>indexed_schema</resource>
  </indexableDocType>
</extension>
```

In this example, the Search service will index all schemas for documents of type `DocType1`, all schemas except `unwanted_schema` for type `DocType2` and only `indexed_schema` for type `DocType3`.

Each ot these indexable schemas will be processed according to the corresponding indexable schema resource if available in the configuration, or to the default one .

In particular, the behavior of data from a given schema is homogeneous across all document types.

11.2.3. The `resource` extension point

This extension point is used to declare an indexable resource. In 5.1M2, the only provided indexable resource type is the `schema` resource type. but the logic will stay the same for future types. Recall that resources are made of *fields*. Here's an example of schema indexing resource, without fields details.

```
<extension target="org.nuxeo.ecm.core.search.service.SearchServiceImpl"
           point="resource" type="schema">

  <resource name="dublincore" prefix="dc" indexAllFields="true">
    <excludedField>issued</excludedField>
    <field name="title" (...) />
  </resource>

  <resource name="book" prefix="bk" type="schema">
    <field name="barcode" (...) />
  </resource>

</extension>
```

The `type` specified that the resource is a docuemnt schema resource.

The `name` and `prefix` attributes are mandatory. They should match the ones from the `schema` extension point of `org.nuxeo.ecm.core.schema.TypeService` as in, e.g.:

```
<schema name="dublincore" prefix="dc" src="schema/dublincore.xsd"/>
<schema name="common" src="schema/common.xsd" />
```

A missing prefix in the core configuration as in the `common` schema declaration in the example above will default to the schema name.

The prefix is important, since all subsequent references to the fields (in queries and raw search results) take the `prefix:fieldName` form.

11.2.4. Field configuration

Fields behavior is designed to be uniform across the different kinds of resources.

11.2.4.1. The `type` attribute

The field type tells the search engine how to process the field. This is a mandatory, *case-insensitive*, attribute.

The following table summarizes the available types. The listed Java classes are guaranteed to work. The backend might implement more converting facilities.

Type	Java classes	Comment
Keyword	String	Meant for technical non binary strings such as vocabulary keys, user identifiers etc. Equality matches are guaranteed to be exact.
Text	String	Upon indexation, these fields are tokenized and analyzed to support fulltext queries.
Path	org.nuxeo.common.utils.Path, String	Dedicated to `STARTSWITH` queries. Equality queries are not guaranteed.
Date	java.util.Calendar	For date and time indexing
Int	int	
Builtin	Reserved for internal use	

11.2.5. Text fields and analyzers

At indexing time, the contents of text fields goes through the process of tokenization and analysis, whose main goal is to provide fulltext search capabilities, usually at the expense of exact matches.

Tokenization means converting the textual content in a stream of words (tokens). During the analysis step, this raw stream is altered to better suit the needs of fulltext searches. It is for example common practice to strip the stream of so-called stop words (most commons words in the language) and to degrade accented characters. One can apply further linguistical processing, such as stemming or synonym analysis.

The name of analyzer to use on a given text field is specified through the `analyzer` attribute. The Search Service acts as a pure forwarder, sending the raw text to the backend, along with the specified analyzer name.

The default value for the `analyzer` attribute is `default`. The attribute is simply ignored for other field types.

11.2.6. Boolean attributes

- To enable queries on a given field, one must set the `indexed` attribute to `true`.
- To make it possible to provide the full field value in search results, one must set the `stored` attribute to `true`. One must keep in mind that the purpose is to present the user a limited yet sufficient set of information along with search results, i.e, in a swifter way than having to fetch it from the core, and not to duplicate the content in the search databases.
- Multivalued fields have to be flagged by setting the `multiple` attribute to `true`.
- Depending on field types and on the processing that's been done by the backend, the possibility to use the field value as a sort key might require some additional resources. To force the backend to give this extra effort, set the `sortable` attribute to `true`.
- The binary attribute is used to mark binary fields, e.g, to trigger conversions (not used in 5.1M2).

11.2.7. Schema resources and fields without configuration

11.2.8. Schema resources

For a schema resource that isn't explicitely declared and nevertheless used, for instance because of an `indexAllSchemas` statement, a default configuration is inferred, with the prefix read from the Core configuration, and fields as below, i.e., as if there was an `indexAllFields="true"` attribute.

11.2.9. Automatic fields configuration

- auto-configured fields are unstored
- If there is only one relevant type from the table above, it is applied.
- The `multiple` attribute is properly set.
- auto-configured String fields get the `keyword` type.

Example 11.1. Relying on automatic field configuration on all fields but one

```
<resource name="dublincore" prefix="dc" indexAllFields="true" type="schema">
  <field name="title" analyzer="default"
         stored="true" indexed="true" type="Text"
         binary="false" sortable="true"/>
</resource>
```

11.3. Programmatic Searching

The search service exposes the `searchQuery` method as a unique entry point. The method takes as input an instance of `ComposedNXQuery`, which encapsulates the parsed NXQL query and a `SearchPrincipal` instance that will be checked against the security indexes, and paging information for the results.

Although the input of `searchQuery` is an already parsed NXQL statement, we'll use NXQL query strings in the sequel for clarity. NXQL query strings are parsed by the method `parse` of the static `org.nuxeo.ecm.core.query.sql.SQLQueryParser` class.

11.3.1. Fields and literals

Within NXQL requests, references to fields values have to follow the "prefix:fieldName" scheme, where prefix and fieldName have been specified through the configuration extension points (recall that for automatically indexed schema resources, the prefix defaults to the one defined in the schema definition).

Literals (constants) follow the JCR specifications. Notably:

- String literals have to be enclosed in single quotes
- Lists have to be enclosed in parenthesizes

Example 11.2. Sample NXQL queries

Recall that "dc" is the prefix for the "dublincore" schema.

```
SELECT * FROM Document WHERE dc:title='Nuxeo book' ORDER BY dc:created DESC
```

11.3.2. `WHERE` statements

Most `WHERE` statements behave as described in the JCR specification, which is itself based on general SQL. Instead of covering every aspect of NXQL, in the current state of this documentation, we'll focus on differences and behaviours that might appear to be counter intuitive.

11.3.2.1. Text fields

Although the Search Service is meant to provide an unified abstraction on the tasks of indexing and querying, text fields have to be somewhat an exception. Indeed, search engines have very different capabilities, depending on provided analyzers. They are nonetheless all expected to provide a direct syntax for full text searches, that an end user can use from, e.g., an input box on a web page. Given the very special kind of constraint that indexing a text field represents, it's not guaranteed that exact matches are supported.

See Section 11.4.4, “Text fields behavior” from the documentation of the Compass backend to get a more concrete view on this (with examples).

11.3.2.1.1. Conclusions

- The backend uses the closest thing to exact matches it supports to treat = predicates.
- The syntax of `LIKE` predicates is backend dependent. It follows the backend's full text query syntax
- The `CONTAINS` from JCR is *not* supported. Use a `LIKE` statement on the main full-text field (see Section 11.3.2.3, “Pseudo fields”).

11.3.2.2. Multi-valued fields

On a multi-valued field, the = operator is true if the right operand belongs to the set of field values. The `IN` operator is true if the intersection between the set of field values and the right operand is non empty.

Example 11.3. Finding documents on which user sally contributed

```
SELECT * FROM Document WHERE dc:contributors = 'sally'
```

Example 11.4. Finding documents on which user sally or harry contributed

```
SELECT * FROM Document WHERE dc:contributors IN ('sally', 'harry')
```

This behavior is in conformance with the JCR specification, which states it the following more general terms:

> In the WHERE clause the comparison operators function the same way they do in XPath when applied to multi-value properties: if the predicate is true of at least one value of a multi-value property then it is true for the property as whole.

11.3.2.3. Pseudo fields

The following fields are available on all document resources. They don't correspond to document fields and aren't configurable, that's why they're called *pseudo-fields*.

The names of these fields are synchronized with constants from the class `BuiltinDocumentFields`. Any use from java code should rely on these.

Constant	Field name	Description
FIELD_FULLTEXT	ecm:fulltext	The default full-text aggregator (string)
FIELD_DOC_PATH	ecm:path	The document path (string)
FIELD_DOC_NAME	ecm:name	The document name (last component of the path, string)
FIELD_DOC_URL	ecm:url	The document URL (string)
FIELD_DOC_REF	ecm:id	The `DocumentRef` as fetched from the core
FIELD_DOC_PARENT_REF	ecm:parentId	The parent `DocumentRef`
FIELD_DOC_TYPE	ecm:primaryType	The Core type (string)
FIELD_DOC_FACETS	ecm:mixinType	The facets (multiple)
FIELD_DOC_LIFE_CYCLE	ecm:currentLifeCycleState	The document life cycle (string)
FIELD_DOC_VERSION_LABEL	ecm:versionLabel	The version label (string)
FIELD_DOC_IS_CHECKED_IN_VERSION	ecm:isCheckedInVersion	A boolean (0 or 1) that states if document is a frozen version (not live nor proxy)

FIELD_DOC_IS_PROXY	ecm:isProxy	A boolean (0 or 1) that states if document is a proxy (targetting other documents)
FIELD_DOC_REPOSITORY_NAME	ecm:repositoryName	The document repository name (string)

11.4. The Compass plugin

Compass is a popular wrapper around Apache Lucene. This plugin allows to use it as a backend for the Search Service.

11.4.1. Configuring Compass

Compass configuration is split in a master XML configuration file and one or several mapping files. The latter specifies the treatments that resources and fields (properties in Compass terminology), while the former is to be used to tune JTA transactions, data sources, and to register configuration of analyzers and converters.

The contents of these files are covered in great details in the Compass 1.1 documentation. In the present documentation, we'll focus on integration matters with the Nuxeo Search Service.

11.4.1.1. Configuration files location

All Compass specific configuration files are relative to the classpath of the compass plugin. A default configuration is provided for Nuxeo EP 5 WebApp. To customize it, one sadly has to put the configuration at the right place within the backend's JAR.

Here is an ant fragment to perform this in a JBoss context, assuming that the Search Service has already been installed in the application server and that the server's deployment directory is stored in the `deploy.dir` property.

```
<copy todir="${deploy.dir}/nuxeo.ear/system/nuxeo-platform-search-compass-plugin-1.0.0-SNAPSHOT.jar"
      overwrite="true" failonerror="false">
  <fileset dir="src/main/resources">
    <include name="myfile.xml" />
  </fileset>
</copy>
```

The backend's JAR is included as a directory in `nuxeo.ear` during the Maven build of `nuxeo-platform-ear` for this single purpose. This is prone to change in the future.

11.4.1.2. Specifying the master configuration XML file name

The Compass backend itself is registered against the Search Service through the `searchEngineBackend` extension point of `org.nuxeo.ecm.core.search.service.SearchServiceImpl`. Your component can use the `configurationFileName` element to specify a path to the master configuration file, like this:

```
<searchEngineBackend name="compass" default="true"
    class="org.nuxeo.ecm.core.search.backend.compass.CompassBackend">
  <configurationFileName>/mycompass.cfg.xml</configurationFileName>
</searchEngineBackend>
```

The default path is `/compass.cfg.xml`.

11.4.2. Global configuration

11.4.2.1. Storage

Compass supports several storage possibilities, called *connections* in Compass configuration objects. The configuration is done trough a Nuxeo Runtime extension point and possibly within the Compass master configuration file. *The extension point always takes precedence over the Compass file*, but can be used to fall back to Compass file, that offers currently more possibilites.

The target is `org.nuxeo.ecm.core.search.backend.compass.CompassBackend`, and the point is `connection`. Contributions are made of a single XML element; the latest one wins.

To set the connection to a file system Lucene store, put a `file` element in the contribution, and set the `path` attribute to the target location. If the path doesn't start with `/`, it will be interpreted as being relative to Nuxeo Runtime's home directory, e.g, `/opt/jboss/server/default/data/NXRuntime` in the default Nuxeo EP installation on JBoss.

Other connection types, notably JDBC, are defined by the Compass configuration file, one has to put the `default` XML element in the contribution, like this:

```
<extension target="org.nuxeo.ecm.core.search.backend.compass.CompassBackend"
           point="connection">
  <default/>
</extension>
```

The default connection is a relative file system one, hosted in the `nxsearch-compass` sub directory of Nuxeo Runtime's home.

11.4.2.2. Analyzers

The master configuration file holds the definition and configuration of analyzers: a lookup name gets associated to an analyzer type and options. The Compass backend makes Compass use directly the name declared in the Search Service as the lookup name, configuration, therefore one has to ensure here that all of these do exist in the Compass configuration.

Together with the registration itself comes the configuration of analyzers. For instance, an analyzer discarding stop words might be given the full list of stop words.

Compass comes with a two predefined analyzers: `default` and `search`. You can reconfigure them as well.

See the relevant part in Compass documentation for details and sample configurations.

11.4.2.3. Converters

Lucene only knows about character strings. Therefore, typed data such as dates and integers must be converted in strings to get into Lucene and back. Compass provides helpers for this and the Compass backend uses them directly.

In the master configuration file, one register available converters in the form of a lookup name and a Java class. Lots of converters are already registered by default, covering most basic types. The `compass.cfg.xml` file shipping with the Compass backend redefines one (the `date` converter).

11.4.3. Mappings for Nuxeo

For the time being, a part of the Search Service configuration has to be duplicated in the Compass mappings XML file.

11.4.3.1. What to describe and syntax

Currently, the Compass backend can't force Compass to use a given converter and/or analyzer on a given field. It *must* therefore be specified in the mappings file, which is itself loaded from the master configuration file. The default name for this file is `nxdocument.cpm.xml`.

Here's a sample, inspired from the mappings file provided with the backend.

```
<?xml version="1.0"?>
<!DOCTYPE compass-core-mapping PUBLIC
    "-//Compass/Compass Core Mapping DTD 1.0//EN"
    "http://www.opensymphony.com/compass/dtd/compass-core-mapping.dtd">

<compass-core-mapping>

  <resource alias="nxdoc"
            sub-index="nxdocs"
            analyzer="default"
            all="false">
    <resource-id name="nxdoc_id"/>

    <resource-property name="dc:created" converter="date"
                       store="yes" index="un_tokenized" />
    <resource-property name="dc:title" analyzer="french" />

  </resource>
</compass-core-mapping>
```

In Compass' terminology, fields are called *properties*. The name of the Compass property corresponding to a Nuxeo indexed field coincides with the field's prefixed name.

Some important remarks:

- Start from the current mappings file shipping with your version of the compass backend and keep it up to date afterwards
- Don't change anything besides `resource-property` elements.
- An exception to the above rules: you may experiment with the `sub-index` attribute, according to your performance needs.

11.4.3.2. Installing the mappings file

Follow the instructions from Section 11.4.1.1, “Configuration files location”

11.4.4. Text fields behavior

Everything in this part applies to fields that have explicitly been declared with a "text" type through the extension point. Anything that's meant for text fields in the mappings configuration files will be ignored if the field has another type.

11.4.4.1. Indexing

At indexing time, the text field is analyzed using the analyzer from the Compass configuration file regardless what has been configured through the Search Service extension point.

11.4.4.2. Searching

Equality statements in WHERE clauses are transformed into the closest thing that Lucene can provide on an analyzed field, namely a phrase query.

On the other hand, LIKE clauses are directly fed to Lucene's `QueryParser`. To search documents whose title starts with "nux", one may write

```
SELECT * FROM Document WHERE dc:title LIKE 'nux*'
```

The following two statements are equivalent. The second one is the QueryParser syntax for phrase queries.

```
... WHERE dc:title='Nuxeo Book'
... WHERE dc:title LIKE '"Nuxeo Book"'
```

Lucene's `QueryParser` syntax is really powerful, you can specify how close two words can be, apply fine grained boosts for the relevance ordering, and more. The only restriction you have on LIKE statements for text

fields within the Compass backend is the choice of field. In other words, the colon character is escaped.

You would need to query date fields, like creation, modification or expiration date, that are provided by Nuxeo Platform. In these cases, it would be interesting to use BETWEEN clauses, associated with DATE keyword, which allows to convert strings as date values.

Example 11.5. Date queries

Documents created in the first term of 2008:

```
... WHERE dc:created BETWEEN DATE '2008-01-01' AND DATE '2008-03-31'
```

Documents modified in may 2007:

```
... WHERE dc:modified BETWEEN DATE '2007-05-01' AND DATE '2007-05-31'
```

You should be aware of the following *trap* in Lucene queries: purely negative queries don't match anything, even if they are themselves nested in a boolean query that has positive statements. The Compass backend uses the standard way to circumvent this limitation, provided that the negative aspect can be seen from the NXQL structure, i.e., not enclosed in a Lucene QueryParser literal.

Example 11.6. Negative queries

Queries that won't return anything:

```
... WHERE dc:title LIKE '-book'
... WHERE dc:title LIKE 'nuxeo' AND dc:title LIKE '-book'
```

Queries that should work as intended (the last three being equivalent):

```
... WHERE dc:title NOT LIKE 'book'
... WHERE dc:title LIKE 'nuxeo' AND dc:title NOT LIKE 'book'
... WHERE dc:title LIKE 'nuxeo' AND NOT dc:title LIKE 'book'
... WHERE dc:title LIKE '+nuxeo -book'
```

11.5. Building a search UI with QueryModel

Chapter 12. Look and feel

12.1. Introduction

The Nuxeo theme notion is wider than the notion attached to the same word in portal concepts. Indeed, the Nuxeo Theme defines all the look and feel of your webapp: composition, layout and graphical appearance. Nuxeo does not aim at developing a portal, i.e. a JSR 168 container, but it authorizes a kind of page and widget management to get some flexibility in the design you want to give to your webapp. The tool to enable you manage those aspects is "Nuxeo Theme editor", that you can make appear with the following command when you are in Nuxeo Web app:

To switch to Nuxeo Theme editor click on the 'Themes Management' link in the user services panel.

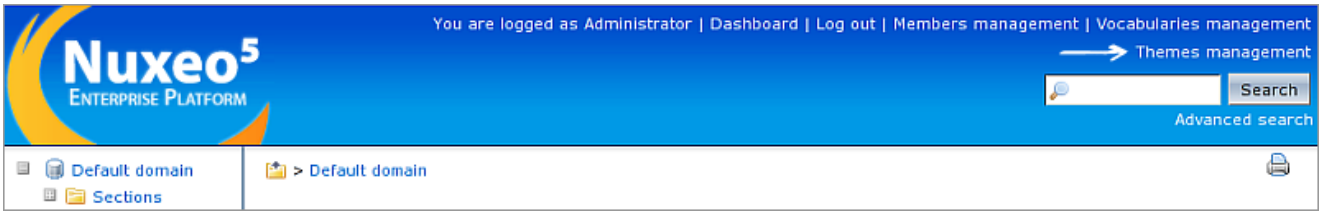

Alternatively, you can type 'shift'+'alt'+'t'. To switch to Nuxeo Theme editor with Mozilla / Firefox < 2.0 type 'alt'+'t'

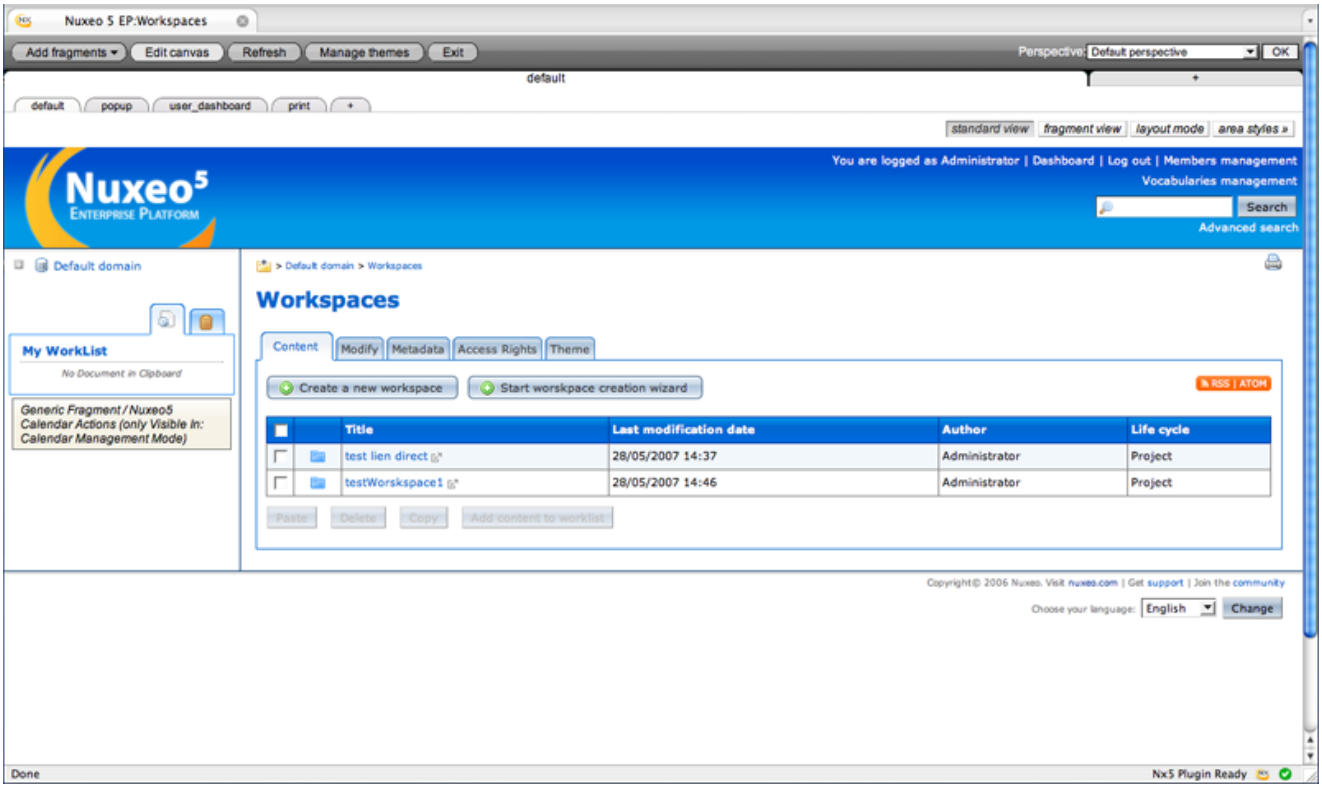

12.2. Principle

One special fragment is the Facelet region: in the properties tab of the editor you can specify the name of a faces to directly integrate it into your page.

To use Nuxeo theme editor, you need to understand its model. The main entity is the theme. Then a theme may have many pages. For each pages, you define a layout (a canvas) and add a list of fragments (widget). The graphical editor uses a tab (theme) and sub-tab (page) system. When you want to add a new page or theme click on "+" at the end of the tab list. For a page you have three possible views:

- **wysiwyg:** you can move the widgets and evaluate the rendering.
- **fragments:** to put the widgets in their placeholder (= a cell). You can put many widgets in a place holder.
- **layout:** you can create new rows and divide the rows into cells, and specify the width of each cell (just edit the nn% on the screen).

You can click on a fragment to edit it. When you edit a fragment you have a multitab editor to specify:

- **custom properties** of the fragment: e.g. the text if it is a text fragment,
- **the graphical object** (the widget) associated
- **the style** of each HTML object that composes the widget.
- then you specify the **perspective** in which the fragment can be seen.

One special fragment is the Facelet region: in the properties tab of the editor you can specify the name of a faces to directly integrate it into your page.

12.3. Mechanism

Nuxeo Editor is done for not having to understand the underground mechanism. Yet it can be good to understand the background to better leverage the tool and its possibilities. The Nuxeo component that manages the customization and extension of Nuxeo EP 5 look and feel is: `org.nuxeo.theme.services.ThemeService`. To register a whole theme (widget, style, layout etc. ...) you need to contribute to the extension point "theme" this way:

```
<extension
    target="org.nuxeo.theme.services.ThemeService" point="themes">
    <theme> <src>META-INF/nxthemes-setup.xml</src>
    </theme> </extension>
```

In the trunk, the default theme is in the webapp project. Having a look inside enables us to discover the main features.

12.3.1. The Elements

The file starts with the elements declaration, we define the pages, the rows (section markup), the cells and the fragments in the cells.

```
<theme name="default"> <layout>
      <page name="default"> <section> <cell> <fragment
      type="generic fragment" /> </cell> <cell> <fragment
      type="generic fragment" /> </cell> <cell> <fragment
      type="generic fragment" /> </cell> </section>
      </page> <page name = ...> . . . </page>
      </layout> </theme>
```

All this markup refers to an `Element` subtype in the java model: `PageElement`, `CellElement`... The **fragment element**, the one that gives the widget oriented capacity to Nuxeo is typed: we have in the default distribution "generic fragment", "action fragments". A typed fragment returns a model to be displayed and edit in the edit mode of the fragments. This model is often (but not always) what we can see in the "properties" tab of the fragment editor. For now there is in the default Nuxeo distribution:

- **generic fragment** (empty model)
- **textual fragment**
- **region fragment**
- **action fragment**

The fragment can also receive a perspective attribute that specifies in which perspective it will be displayed (the fourth tab in the fragment editor). You can then propose to Nuxeo user the same kind of experience you have in Eclipse. The perspective are specified in the **perspective** extension point of `!ThemeService` component.

12.3.2. The format

Then, once you declared all the elements, you can format them through different axes :

- their layout
- their rendering (their view)
- their style

To do this, you put, inside the `<theme>` markup, children markup from those types:

- layout --> `<layout element = ...>`
- rendering --> `<widget element ...>`
- style --> `<style element = ...>`

Those 3 markups use the attribute element to get the reference to which element of your skeleton they will be applied:

```
element="page[3]/section[3]/cell[1]"
```

12.3.2.1. The rendered widget

Indeed each element is rendered by a view.

```
<widget element="page"> <view>page
        frame</view> </widget>
```

This view is defined like this (with another extension point of `!ThemeService: views`):

```
<view
        name="page frame"> <element-type>page</element-type>
        <format-type>widget</format-type>
        <class>org.nuxeo.theme.jsf.views.JSFView</class>
        <template>nxthemes/jsf/widgets/page-frame.xml</template>
        </view>
```

We can see that a view is associated to an element type. The element types contributions should be reserved to Nuxeo only (one should manage with existing ones). The template markup gives the html/faces/text code to be used for rendering the view. Notice that in the view, you can access the fragment model data through the EL call `nxthemesInfo.model`.

```
<h:outputText escape="false"
          value="#{nxthemesInfo.model.body}" />
```

12.3.2.2. The layout

The layout properties are given like this (still under the `<theme>` markup) :

```
<layout element="page[3]/section[3]/cell[1]">
        <width>50%</width> <padding>20px</padding>
        <margin>0</margin> </layout>
```

This enables to adjust the position of the fragments inside.

12.3.2.3. The style

Then comes the `style` properties. Again, you apply them to an element:

```
<style
        element="page[1]/section[3]/cell[1]|page[3]/section[4]/cell[1]">
        <selector path=""> <color>#757575</color>
        <border-style>solid none none none</border-style>
        <border-color>#003366</border-color>
        <border-width>1px</border-width> <background> #FFF
        url(/nuxeo/img/gray_gradient.gif) top left repeat-x
        </background> <padding>5px 15px 5px 5px</padding>
        </selector> <selector path="div">
        <font-size>9px</font-size> </selector>
        </style>
```

The selector specifies the markup to which the defined style is applied. The style definition used is the one of the deeper upper-element that has a style definition specified. To be exhaustive, we need to present the filter system (TODO)

12.3.3. The negotiation

We have seen how to define a theme. Now we need to see how a theme is applied. More precisely, how do I manage the choice of the page I will display? In fact, the Nuxeo Theme framework proposes many ways to specify the theme applied to the webapp for a given view:

- with a cookie
- with a request parameter (?theme= ...)
- with an association between a JSF view id and a theme

So how to manage priority when more than one parameter is passed to the framework? The `!ThemeService` component has another extension point to achieve this: the `negotiation` extension point. Not only can it be used to select a theme but it also works with other objects as we will see later.

```
<negotiation object="theme"
      strategy="nuxeo5">
      <scheme>org.nuxeo.theme.jsf.negotiation.theme.RequestParameter</scheme>
      <scheme>org.nuxeo.theme.jsf.negotiation.theme.CookieValue</scheme>
      <scheme>org.nuxeo.theme.jsf.negotiation.theme.ViewId</scheme>
      <!-- local theme (specific to nuxeo5) -->
      <scheme>org.nuxeo.ecm.webapp.theme.LocalTheme</scheme>
      <scheme>org.nuxeo.theme.jsf.negotiation.theme.DefaultTheme</scheme>
      </negotiation>
```

As we can see in the example, the negotiation point defines the order in which the different methods for obtaining the current theme information are applied. This negotiation feature also applies to other Nuxeo Theme objects like the engine, the mode, the perspective.

12.3.4. The engine

Engine and filter are two notions that work together. An engine is the combination of different filters, and a filter is a "sub-unit" of rendering. So the engine is the global renderer of your web app. From the elements skeleton, it will generate the graphical appearance, passing each element through different black boxes, depending on the type of the element. Here is the definition of the default engine of Nuxeo.

```
<engine name="default"> <renderer
      element="theme"> <filter>add widget</filter>
      <filter>collect xmlns</filter> </renderer>
      <renderer element="page"> <filter>add widget</filter>
      <filter>set layout</filter> <filter>set
      style</filter> </renderer> <renderer
      element="section"> <filter>add widget</filter>
      <filter>set layout</filter> <filter>set
      style</filter> </renderer> <renderer element="cell">
      <filter>add widget</filter> <filter>set
      layout</filter> <filter>set style</filter>
      </renderer> <renderer element="fragment">
      <filter>control fragment visibility</filter>
```

```
      <filter>add widget</filter> <filter>set
      style</filter> <filter>write fragment tag</filter>
      </renderer> </engine>
```

The engine, that you register in the `ThemeService` component through the "engines" extension point lets you add for each type of element some "filters" that will do some work around the markup content at rendering time. Nuxeo already uses filters like the style fitler, that put the style definition you chose, the layout filter, the "drag'n drop" filter... One interesting use of the filters is illustrated with the Nuxeo Theme editor: when you type **shift+alt+t**, the display changes and you are in the WYSIWYG Nuxeo theme editor. But all the components that make your page are still here, they just look a bit different because, some different filters are used. For instance, because of the drag'n drop filter presence, you can move the fragments.

12.3.5. Resource management

Graphical components may need some external resources such as CSS or JavaScript libraries. Nuxeo theme has an embedded resource management system that at rendering time automatically computes the list of the files needed for rendering a page. Resources are served using gzip compression when supported by the browser. JavaScript resources are also compressed using Dojo's ShrinkSafe. Finally all files of a same type (.css or .js) are concatenated. This reduces the number of individual downloads and it enables to manage dependencies between resources. Indeed, at declaration time you can specify the dependencies for a given resource:

```
<extension
      target="org.nuxeo.theme.services.ThemeService" point="resources">
      <resource name="controls.js">
      <path>nxthemes/jsf/scripts/scriptaculous/controls.js</path>
      <require>effects.js</require>
      <require>prototype.js</require> </resource>
      </extension>
```

Then, when you register the view associated to an element, you specify the resources it needs:

```
<view name="nuxeo5
      clip board"> <format-type>widget</format-type>
      <class>org.nuxeo.theme.jsf.views.JSFView</class>
      <template>incl/user_clipboard.xhtml</template>
      <resource>dragdrop.js</resource>
      </view>
```

12.3.6. Application

The last concept you need to know about to completely control the look and feel of your application is the application extension point:

```
<extension
      target="org.nuxeo.theme.services.ThemeService" point="applications">
      <application root="/nuxeo"> <negotiation>
      <strategy>nuxeo5</strategy>
      <default-engine>default</default-engine>
      <default-theme>default/default</default-theme>
      <default-perspective>default</default-perspective>
      </negotiation> <resource-caching>
      <lifetime>36000</lifetime> </resource-caching>
      <view id="/create_relation_search_document_popup.xhtml">
      <theme>default/popup</theme> </view> <view
      id="/user_dashboard.xhtml">
      <theme>default/user_dashboard</theme> </view> <view
      id="/view_calendar.xhtml">
      <perspective>view_calendar</perspective> </view>
      <view id="/print.xhtml">
      <perspective>print</perspective> </view>
      </application> </extension>
```

As you can see in the example, an **application** is associated to a web-app root context. There you specify the strategy (a negotiation grouping feature), the default engine, the default theme and perspective. You also specify the caching policy and there you also declare the JSF view id / theme association that we went through earlier in this tutorial.

12.4. Customizing the theme

Eventually, all theme and styling work will be done in the Theme Editor. For now, we have to use both the editor and the file `theme-default.xml` in `nuxeo/nuxeo-platform/nuxeo-platform-webapp-core/src/main/resources/META-INF/`.

What can be done in the editor: page layout, widget moving, fragment styling, page/section/cell preset borders and backgrounds

What must be done in `theme-default.xml` : commons styles and their inheritance

In addition to the `theme-default.xml` come palettes: a bunch of presets for colors, backgrounds, fonts and other css attributes. Nuxeo EP 5 supports text palettes and GIMP/Photoshop palettes (for the colors).

When you add images or modify `theme-default.xml`, you have to redeploy your Nuxeo 5.

In case of doubt, try using the editor, because the produced code is much cleaner and compliant than anything you would write manually :-)

12.4.1. Modifying the current theme using `theme-default.xml`

The file `theme-default.xml` is structured as follows:

- Pages and their layout
- widgets in pages
- definition of predefined styles (using preset values from palettes)
- cell and fragments styling

In `theme-contrib.xml` we have our theme called:

```
<!-- themes --> <extension
        target="org.nuxeo.theme.services.ThemeService" point="themes">
        <theme> <src>META-INF/theme-default.xml</src>
        </theme> </extension>
```

In the editor, in Manage Themes tab, it gives:

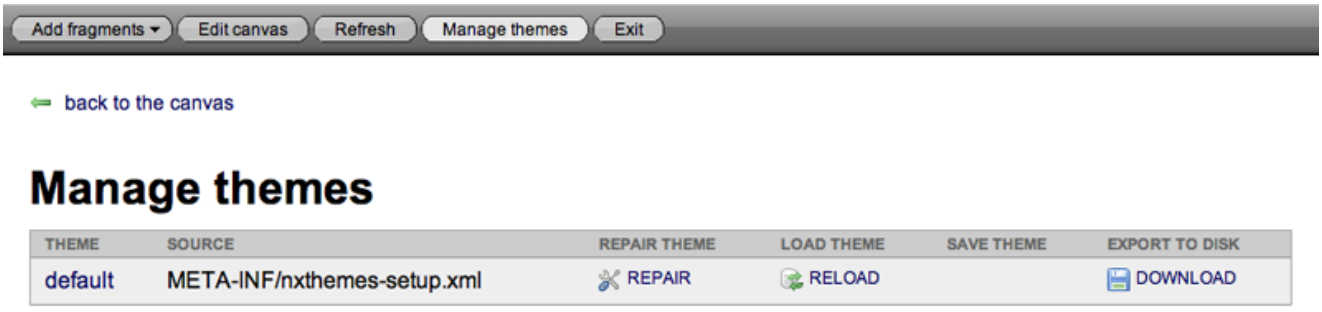

This file is deployed in JBoss. If you modify the theme using the editor all changes will be lost so think of downloading the theme to your Desktop, to replace the `theme-default.xml` in your local copy of Nuxeo EP 5.

A good way of working with this file is to add your working copy in `theme-contrib.xml`. It is possible in NXThemes to load several themes and page.

Add your file(s) in `themes-contrib.xml`, for example:

```
<theme>
        <src>file:///path/to/sources/nuxeo/nuxeo-platform/nuxeo-platform-webapp-core/src/main/resources/META-INF
        </theme>
```

After a redeployment, in the 'Manage Themes' section we now have a theme that can be reloaded directly from the file-system!

Starting from there, here are two ways of developing smartly:

- Edit your `theme-default.xml` in Eclipse (or in your XML editor) then go to theme editor, in Manage Themes tab and click "Reload" on your local file: you can directly see the changes you made in the XML source.
- Modify the theme inside the editor, then go to Manage Themes tab and click on "Save" action. All changes will be saved in the file.

12.4.1.1. Playing with palettes.

The palettes are in
`nuxeo/nuxeo-platform/nuxeo-platform-webapp-core/src/main/resources/themes.palettes/`

They are called in
`nuxeo/nuxeo-platform/nuxeo-platform-webapp-core/src/main/resources/OSGI-INF/theme-contrib.xml`:

```
<!-- Styles presets --> <extension
          target="org.nuxeo.theme.services.ThemeService" point="presets">
          <palette name="Nuxeo default fonts"
          src="themes/palettes/nxfonts.properties" category="font" />
          <palette name="Nuxeo psd colors"
          src="themes/palettes/nxcolors.aco" category="color" />
          <palette name="Nuxeo default backgrounds"
          src="themes/palettes/nxbackgrounds.properties" category="background"
          /> </extension>
```

There are 3 default palettes:

- `nxbackgrounds.properties` that specifies the banner's css background and the shadow under it
- `nxcolors.aco` that contains nuxeo default colors in a photoshop palette format
- `nxfonts.properties` that contains default css font, small and 4 levels of titles

The easiest way for you to customize yout Nuxeo EP 5 app is to modify the existing palettes!

For example, in `nxfonts.properties` change the line

```
default=11px Verdana, Arial,
          sans-serif
```

to

```
default=12px Courier, serif
```

Then all the fonts of the app will be changed to your new value!

We advise you to add your own color palette.

12.4.1.2. Definition of a predefined style

Currently in `nxthemes-setup.xml` we have a style named default buttons, which is defined as:

```
<style name="default buttons"> <selector
          path="input.button">
          <background>url(/nuxeo/img/button_1.gif) 0 0 repeat-x
```

```
          #e3e6ea</background> <font preset="default (Nuxeo default
          fonts)"/> <margin>5px 10px 10px 0px</margin>
          <color>#000</color>
          <border-style>solid</border-style>
          <border-width>1px</border-width>
          <border-color>#ccc #666 #666 #ccc</border-color>
          <padding>2px 5px 2px 5px</padding>
          <cursor>pointer</cursor> </selector> <selector
          path="input.button:hover"> <color preset="white (Nuxeo psd
          colors)"/> <font preset="default (Nuxeo default fonts)"/>
          <background>url(/nuxeo/img/button_2.gif) 0 0 repeat-x
          #3f89ef</background> <border-color>#0099ff #0066cc
          #0066cc #0099ff</border-color>
          <border-style>solid</border-style>
          <border-width>1px</border-width> </selector>
          <selector path="input.button[disabled=disabled],
          input.button[disabled=disabled]:hover">
          <color>#c1c1c1</color> <font preset="default (Nuxeo
          default fonts)"/>
          <background>url(/nuxeo/img/button_disabled.gif) 0 0 repeat-x
          #ebeff4</background> <border-color>#ccc #999 #999
          #ccc</border-color> <cursor>default</cursor>
          <border-style>solid</border-style>
          <border-width>1px</border-width> </selector>
          </style>
```

We can see that:

- a style that does not apply to an `element` is name
- inside this style, several HTML attributes/classes are called
- palette preset are called, such as the font attribute

Predefined styles are also a good way of efficiently changing the look of your application because you need to change the CSS only once!

12.4.1.3. Using a predefined style

Later in the file we notice that the 'user services' fragment takes the default buttons style preset:

```
<!-- user services --> <style
          element="page[1]/section[1]/cell[2]/fragment[1]|page[3]/section[1]/cell[2]/fragment[1]"
          inherit="default buttons">
```

It means that the styles defined for the buttons will be applied to the 'user services' fragment (user links and search in the banner).

12.4.2. Modifying the current theme

12.4.2.1. Styling the theme using the editor

As we explained earlier, layout editing and local styling can be done in the theme editor.

In the editor, click on an element you want to style, click "Edit" in the Menu. Here we chose the RSS/Atom link button

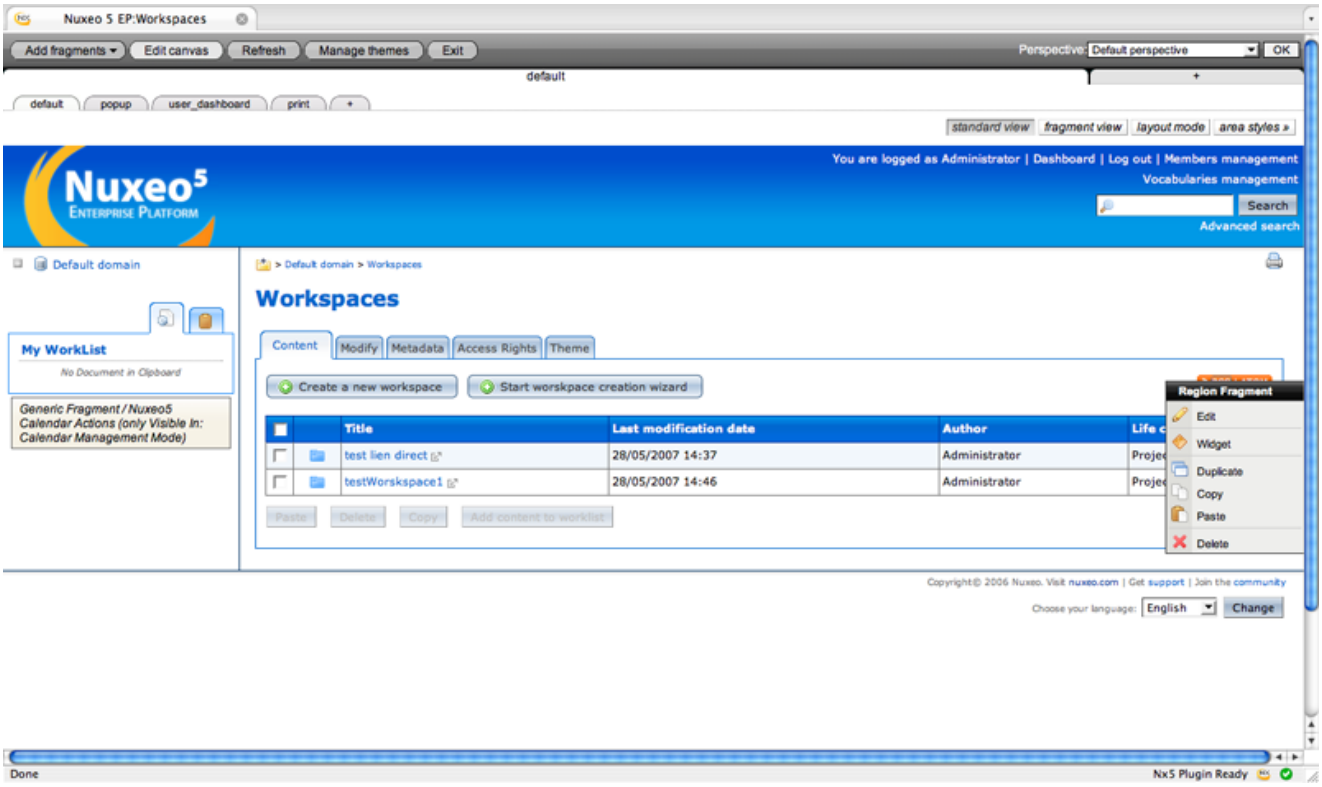

Access the Style tab.

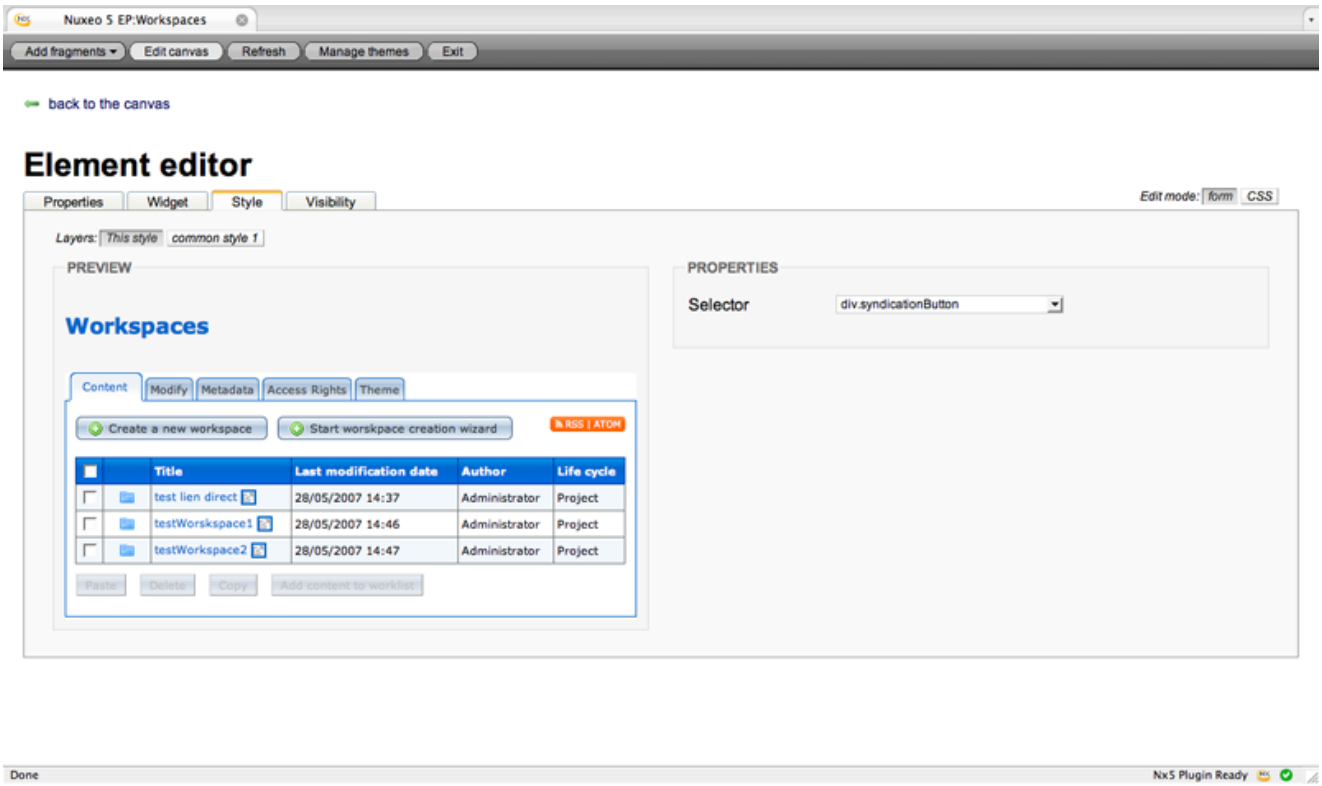

The existing selectors are on the right in the Properties box, otherwise move the mouse over the preview area and click on an element to create a CSS selector path.

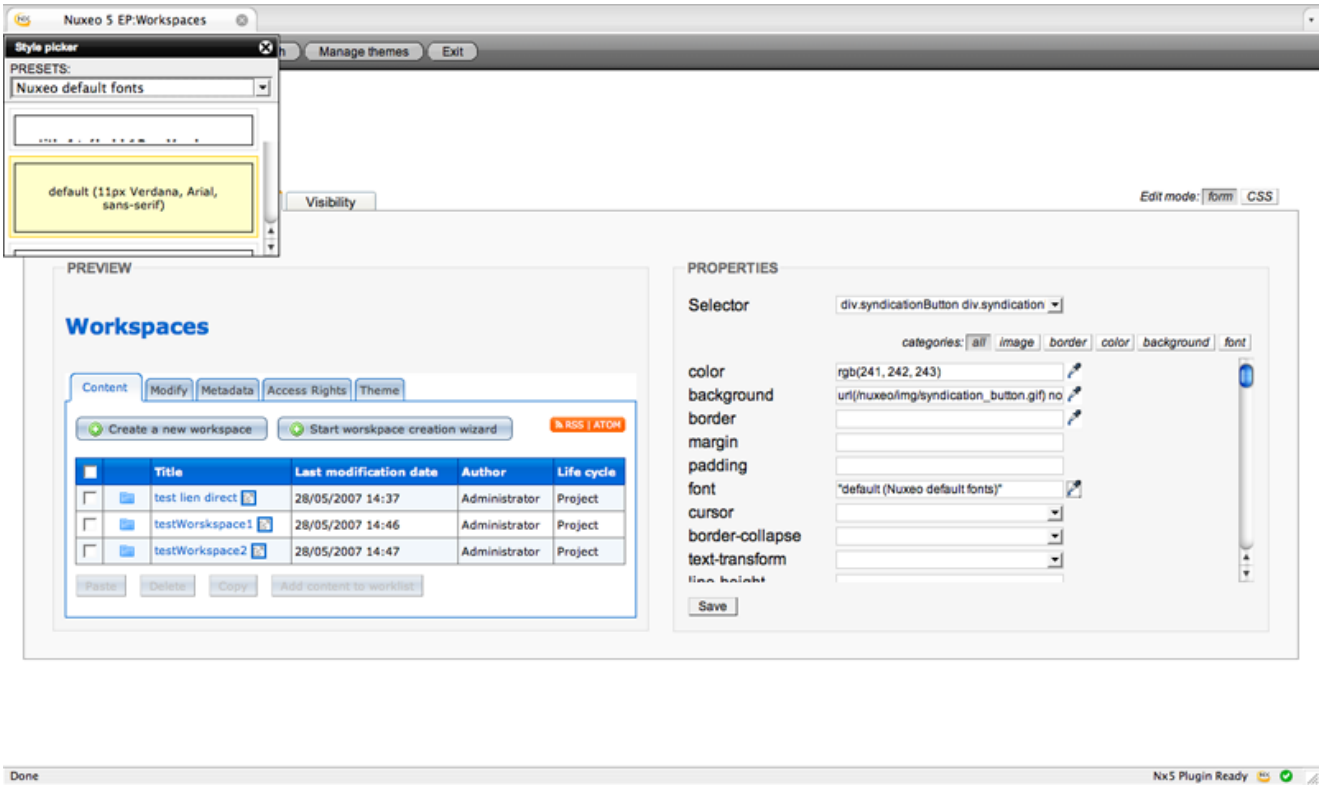

We choose to change the small font preset to the default one. As you see, the Style picker shows all the palettes and all the presets are rendered. We remove the background property for the syndication links button and add a preset background-color, our RSS/Atom button is all changed now:

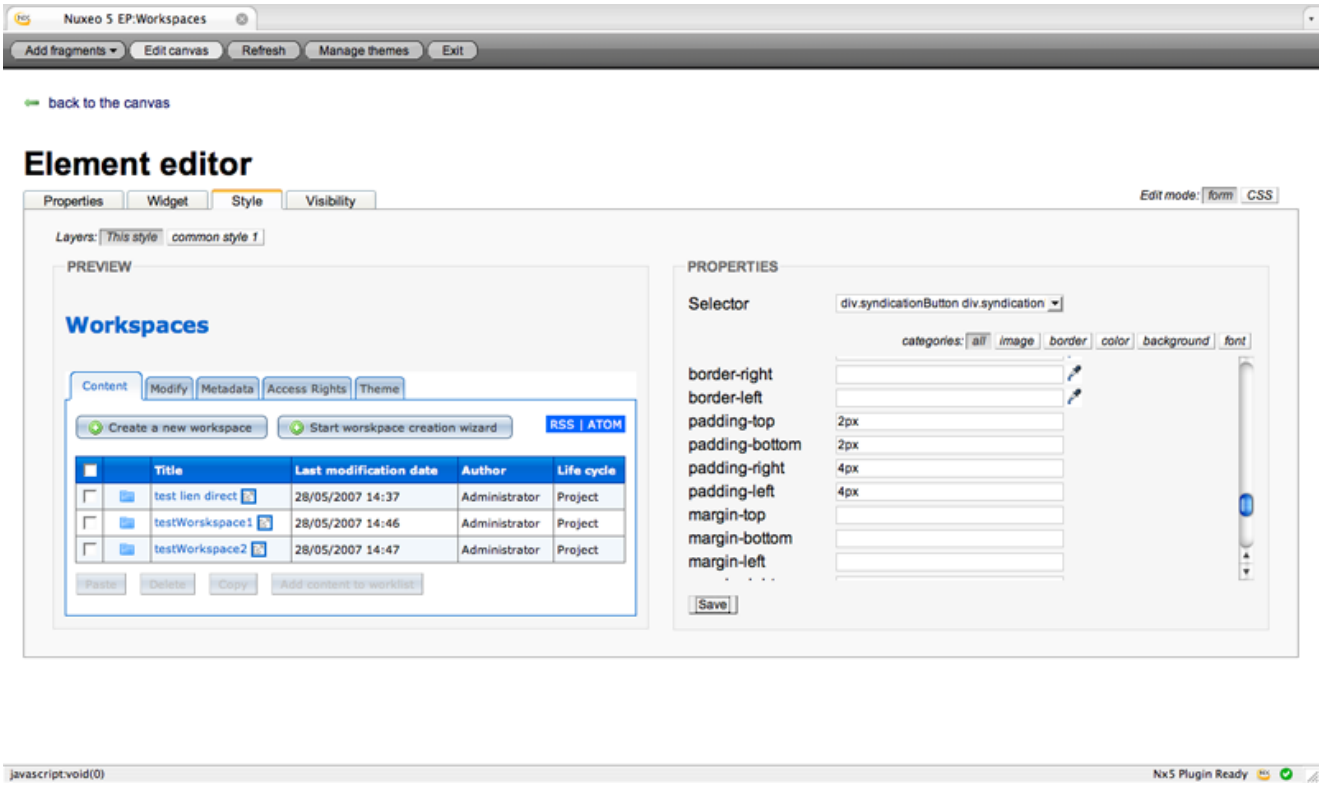

When you are done with managing your theme you might want to save it to your local copy of Nuxeo. Just go in the Manage Themes tab, download the custom theme to your computer, then put it in your repository.

Congratulation, you have just customized the Nuxeo EP theme!

12.4.2.2. Modifying the theme by adding/modifying a fragment

You may want to modify an existing fragment to customize your project, let's say you want your compagny logo instead of Nuxeo EP's and you own corporate links in the footer. We won't create & declare new fragments (as we saw, fragments and their resources are defined in `theme-contrib.xml`), we'll use the default-ones to override Nuxeo EP's, considering you have your own project using Nuxeo EP default as made in the *sample project*.

Here are the steps do to so:

- copy your logo (let's call it `corporate_logo.gif`) to `your.project/src/main/resources/nuxeo.war/img`
- copy and paste `logo.xhtml` and `footer.xhtml` from `nuxeo/nuxeo-platform/nuxeo-platform-webapp/src/main/resources/nuxeo.war/incl` to `your.project/src/main/resources/nuxeo.war/incl` so it's overridden when doing your *ant*.

Tip

This is a general principle for *nuxeo.war* folder. The contents of the */img/* folder of your app are the contents of Nuxeo EP's default `.../nuxeo.war/img` folder. Every specific resource in `your.project/.../nuxeo.war/img` come in addition of what is already in default `.../nuxeo.war/img` if non-existing there with same filename, or come **instead** of what is existing in default `.../nuxeo.war/img` if same filename.

- edit `logo.xhtml` that currently contains

```
<div xmlns:h="http://java.sun.com/jsf/html"
            xmlns:ui="http://java.sun.com/jsf/facelets"
            xmlns:f="http://java.sun.com/jsf/core"
            xmlns:t="http://myfaces.apache.org/tomahawk"
            xmlns:nxu="http://nuxeo.org/nxweb/util"
            xmlns:c="http://java.sun.com/jstl/core" class="menu">
            <div> <h:form> <h:commandLink
            action="#{navigationContext.goHome()}"> <h:graphicImage
            value="#{logoHelper.logoURL}" alt="Nuxeo EP 5" width="194"
            height="99" /> </h:commandLink> </h:form>
            </div> </div>
```

- change the line `<h:graphicImage value="#{logoHelper.logoURL}" alt="Nuxeo EP 5" width="194" height="99" />` for something like `<img src="/nuxeo/img/corporate_logo.gif" alt="My corporate logo" />` and save your changes
- edit `footer.xhtml` that currently contains

```
<div xmlns:h="http://java.sun.com/jsf/html"
              xmlns:ui="http://java.sun.com/jsf/facelets"
              xmlns:f="http://java.sun.com/jsf/core"> <ui:insert
              name="footer"> Copyright
              <f:verbatim>&amp;copy;</f:verbatim> 2006 Nuxeo.
              Visit <!--   --> <h:outputLink
              value="http://www.nuxeo.com"> <h:outputText
              value="nuxeo.com" /> </h:outputLink> | Get <!--
                --><h:outputLink
              value="http://www.nuxeo.com/en/services/support/">
              <h:outputText value="support" /> </h:outputLink> |
              Join the <!--   --> <h:outputLink
              value="http://www.nuxeo.org/sections/community">
              <h:outputText value="community" /> </h:outputLink>
              <br /> <h:form> <h:outputText
              value="#{messages['label.selectLocale']}" />
              <h:selectOneMenu value="#{localeSelector.localeString}"
              styleClass="langSelect"> <f:selectItems
              value="#{localeSelector.supportedLocales}"/>
              </h:selectOneMenu> <h:commandButton
              action="#{localeSelector.select}"
              value="#{messages['command.changeLocale']}" class="langSubmit"
              /> </h:form> </ui:insert>
              </div
```

- Change from `Copyright` to `<br />` by something like `<a href="http://yoursite.com">My Corporate Site</a>`, save your changes

- Do an *ant* on your projet, rerun your *jboss* and appreciate the changes...

Congratulation, you have just customized some Nuxeo EP fragments!

12.4.3. Adding a new theme and its pages

We rarely create a theme from scratch. Since the default, popup and dashboards are used in every projetc, we usually duplicate the default theme and use it as a basis. But let's pretend you want to add a completely new and custom theme to your project. Here are the steps.

Note

We assume you are familiar to Nuxeo EP way and have read the theme section above!

12.4.3.1. Create a new theme

- Log in as *Administrator*
- Go to the *Manage Theme* view
- Click the *plus* tab, right next to the themes names

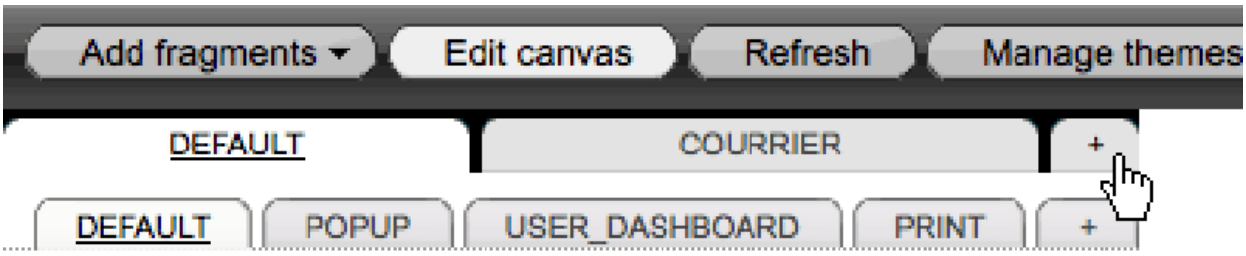

- Define a (smart) name. We are currently using "cuztom"

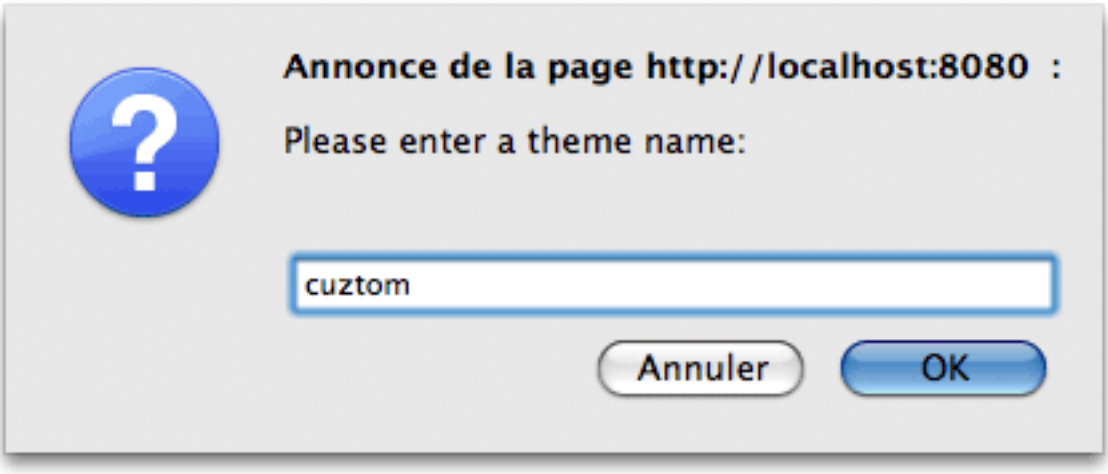

- You have a page called "default", toy with it, add sections and fragments, color the areas so your theme is not empty. We recommend to add a region fragment with name set as *body* so the main content is displayed.

- You can add pages to your theme by clicking the *plus* tab next to the pages names.
- Go the *manage themes* tab
- Download your theme

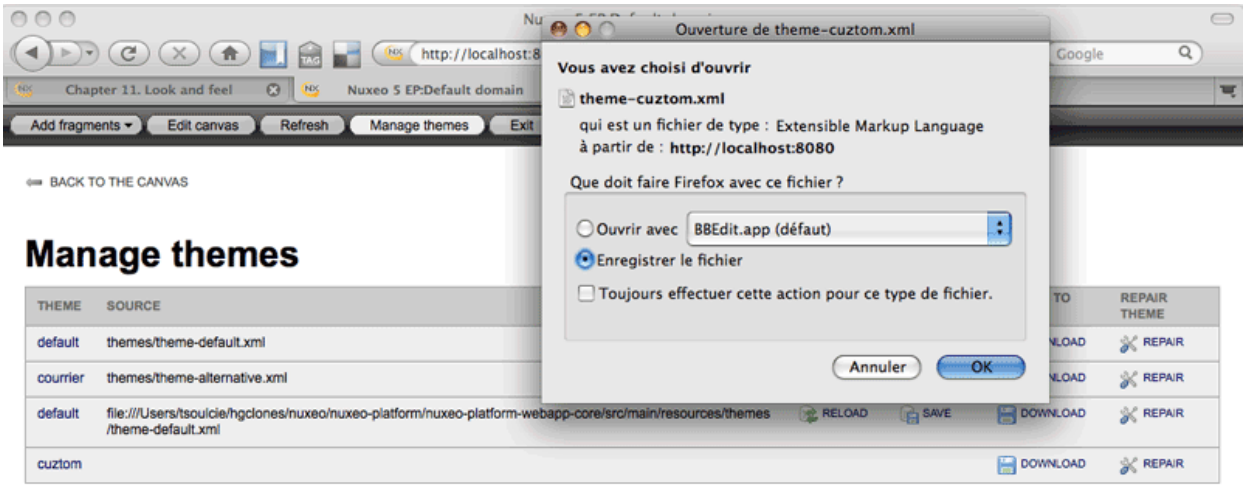

12.4.3.2. Setting the files

You will need a few files/declaration to set your new theme. What we will be overriding in our project is `/nuxeo/nuxeo-platform/nuxeo-platform-webapp-core/src/main/resources/`.

Make sure that:

- You copied your downloaded cuztom theme (or duplicate default nuxeo theme, of course) to `/yourproject/.../resources/themes`
- `META-INF/META-INF.MF` calls your `OSGI-INF/cuztom-theme-contrib.xml` as component
- You have a `OSGI-INF/cuztom-theme-contrib.xml` with inside:
  - **The `component name` set is your project's**
  - The `theme extension point` contains your JBoss's theme file
  - The `theme extension point` contains your local theme file (if you plan to work that way)

theme-contrib.xml | theme-cuztom.xml | Untitled.html | MANIFEST.MF

```
<?xml version="1.0"?>

<component name="com.nuxeo.projects.cuztom.theme.contrib">

  <extension target="org.nuxeo.theme.services.ThemeService" point="themes">

    <theme>
      <src>themes/theme-cuztom.xml</src>
    </theme>

    <theme>
      <src>file:///Users/tsoulcie/cuztomproject/src/main/resources/themes/theme-cuztom.xml</src>
    </theme>

  </extension>
```

- The `applications extension point` set the negociation with your theme as default theme

```
<!-- application settings -->
<extension target="org.nuxeo.theme.services.ThemeService"
  point="applications">

  <application root="/nuxeo">

    <negotiation>
      <strategy>nuxeo5</strategy>
      <default-engine>default</default-engine>
      <default-theme>cuztom/default</default-theme>
      <default-perspective>default</default-perspective>
    </negotiation>
```

- If you have specific fragments or palettes, declare them there.

Your new theme is now part of your project and set as default. Your can *ant* your app and re-run JBoss.

Tip

A good way to name your files is to add your project's name before the current Nuxeo filename. For an example, `theme-contrib.xml` may become `cuztom-theme-contrib.xml`.

12.4.3.3. Working the theme

Your now all set to create your own design with all the tools explained in the sections above:

- Create fragments, declare theme in `cuztom-theme-contrib.xml` and drop them into your pages using the theme editor
- Add palettes, declare them, call them in your theme-file.
- Create generic style and make elements inherit of it.
- Modify the structure of the pages with the theme editor and save your changes in your local files
- Add style in your local files and reload the theme using the Manage Theme tab in theme editor.
- Add images and icon into your nuxeo.war folder and call them in the style, in the actions-contrib.xml or in the fragments.

Have fun and send us your creations!

Chapter 13. Authentication, Users & Groups Management

13.1. Introduction

In Nuxeo EP, the concept of a *user* is needed for two main reasons:

- Users are needed for authentication and authorization to work,
- Users have associated information that can be displayed, for instance to display someone's full name or email address.

An abstraction, the *`UserManager`*, centralizes the way a Nuxeo EP application deals with users (and groups of users). The `UserManager` is queried by the platform's *`LoginModule`* when someone attemps to authenticate against the framework. It is also queried whenever someone wants the last name or email of a user for instance, or to get all users having "Bob" as their first name.

13.2. Users and Groups configuration

The data about users (login, password, name, personal information, etc.) and the groups they belong to (simple members, or any application-related group) are managed through the Directory abstraction. This means that users can be stored in LDAP or in SQL, and groups can be part of the LDAP tree or stored in a SQL table, but the application doesn't see the difference as long as the connectors are configured properly.

13.2.1. Schemas

To define a new source of users, you simply define a new directory (or redefine the default one) to use different connection and schema information. We'll use an example where the pet-name field is added to the user's schema.

Let's start by defining our new schema in a `.xsd` file, `myuser.xsd`:

```
<?xml version="1.0"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
    xmlns:nxs="http://www.nuxeo.org/ecm/schemas/user"
    targetNamespace="http://www.nuxeo.org/ecm/schemas/user">

  <xs:include schemaLocation="base.xsd" />

  <xs:element name="username" type="xs:string" />
  <xs:element name="password" type="xs:string" />
  <xs:element name="email" type="xs:string" />
  <xs:element name="firstName" type="xs:string" />
  <xs:element name="lastName" type="xs:string" />
  <xs:element name="company" type="xs:string" />

  <xs:element name="petName" type="xs:string" />

  <!-- inverse reference -->
  <xs:element name="groups" type="nxs:stringList" />

</xs:schema>
```

This schema must be registered in an extension point:

```
<extension target="org.nuxeo.ecm.core.schema.TypeService" point="schema">
  <schema name="myuser" src="myuser.xsd" />
</extension>
```

TODO: FG groups

13.2.2. Directories

The user schema can now be used when we define a new directory, `MyUserDirectory`. A SQL directory is defined like this:

```
<extension target="org.nuxeo.ecm.directory.sql.SQLDirectoryFactory" point="directories">
  <directory name="MyUserDirectory">

    <schema>myuser</schema>
    <idField>username</idField>
    <passwordField>password</passwordField>

    <dataSource>java:/nxsqldirectory</dataSource>
    <table>myusers</table>
    <dataFile>myusers.csv</dataFile>
    <createTablePolicy>on_missing_columns</createTablePolicy>

    <references>
      <inverseReference field="groups" directory="groupDirectory"
        dualReferenceField="members" />
    </references>

  </directory>
</extension>
```

And we can provide a file, `myusers.csv`, which will be used to populate the table if it is missing:

```
username, password, firstName, lastName, company, email, petName
bob,bobSecret,Bob,Doe,ACME,bob@example.com,Lassie
```

If instead we had used an LDAP directory, the configuration would look like:

```
<extension target="org.nuxeo.ecm.directory.ldap.LDAPDirectoryFactory" point="servers">
  <server name="default">
    <ldapUrl>ldap://localhost:389</ldapUrl>
    <bindDn>cn=manager,dc=example,dc=com</bindDn>
    <bindPassword>secret</bindPassword>
  </server>
</extension>

<extension target="org.nuxeo.ecm.directory.ldap.LDAPDirectoryFactory" point="directories">
  <directory name="MyUserDirectory">

    <schema>myuser</schema>
    <idField>username</idField>
    <passwordField>password</passwordField>

    <server>default</server>
    <searchBaseDn>ou=people,dc=example,dc=com</searchBaseDn>
    <searchClass>inetOrgPerson</searchClass>
    <searchScope>subtree</searchScope>

    <fieldMapping name="username">uid</fieldMapping>
    <fieldMapping name="password">userPassword</fieldMapping>
    <fieldMapping name="email">mail</fieldMapping>
    <fieldMapping name="firstName">givenName</fieldMapping>
    <fieldMapping name="lastName">sn</fieldMapping>
    <fieldMapping name="company">o</fieldMapping>

    <references>
      <inverseReference field="groups" directory="groupDirectory"
        dualReferenceField="members" />
    </references>

  </directory>
</extension>
```

Detailed configuration on SQL Directories and LDAP Directories can be found in Chapter 18, *Directories and Vocabularies*.

13.2.3. UserManager

We can now tell the UserManager that this directory should be the one to use when dealing with users:

```
<extension target="org.nuxeo.ecm.platform.usermanager.UserService" point="userManager">
  <userManager>
```

```
    <users>
      <directory>MyUserDirectory</directory>
      <emailField>email</emailField>
      <searchFields append="true">
        <searchField>username</searchField>
        <searchField>firstName</searchField>
        <searchField>lastName</searchField>
        <searchField>myfield</searchField>
      </searchFields>
    </users>

  </userManager>
</extension>
```

This configuration also sets the email field, and search fields that have to be queried when searching for users. It can be completed to set the anonmymous user, add virtual users, or set the group directory properties.

```
<extension target="org.nuxeo.ecm.platform.usermanager.UserService" point="userManager">
  <userManager>

    <users>

      <directory>MyUserDirectory</directory>
      <emailField>email</emailField>
      <searchFields append="true">
        <searchField>username</searchField>
        <searchField>firstName</searchField>
        <searchField>lastName</searchField>
        <searchField>myfield</searchField>
      </searchFields>
      <listingMode>tabbed</listingMode>

      <anonymousUser id="Anonymous">
        <property name="firstName">Anonymous</property>
        <property name="lastName">User</property>
      </anonymousUser>
      <virtualUser id="MyCustomAdministrator" searchable="false">
        <password>secret</password>
        <property name="firstName">My Custom</property>
        <property name="lastName">Administrator</property>
        <group>administrators</group>
      </virtualUser>
      <virtualUser id="MyCustomMember" searchable="false">
        <password>secret</password>
        <property name="firstName">My Custom</property>
        <property name="lastName">Member</property>
        <group>members</group>
        <group>othergroup</group>
        <propertyList name="listprop">
          <value>item1</value>
          <value>item2</value>
        </propertyList>
      </virtualUser>
      <virtualUser id="ExistingVirtualUser" remove="true" />

    </users>

    <defaultAdministratorId>Administrator</defaultAdministratorId>
    <userSortField>lastName</userSortField>
    <userPasswordPattern>^[a-zA-Z0-9]{5,}$</userPasswordPattern>

    <groups>
      <directory>somegroupdir</directory>
      <membersField>members</membersField>
      <subGroupsField>subgroups</subGroupsField>
      <parentGroupsField>parentgroup</parentGroupsField>
      <listingMode>search_only</listingMode>
    </groups>
    <defaultGroup>members</defaultGroup>
    <groupSortField>groupname</groupSortField>

  </userManager>
</extension>
```

The anonymous user represents a special kind of virtual user, used to reprensent users that do not need to log in the application. This feature is used in conjunction with the anonymous plugin (see next chapter).

Virtual users can be added for authentication. Properties are used to create the appropriate model as if user was retrieved from the user directory. This is a convenient way to add custom users to the application when the user directory (using LDAP for instance) cannot be modified. Virtual users with the "administrators" group will have the same rights than the default administrator.

The default administrator id can be set either to an existing or virtual user.

The group directory can also be configured to define the groups hierarchy and the contained users. This configuration has to match the user directory inverse references.

Every authenticated user will be placed in the configured default group. This group does not need to exist in the backing group directory, nor does any other group listed in virtual users configuration.

13.2.4. User Management Interface

Finally, because the different fields of a user are visible when creating or viewing them through the user management screens, we redefine the layout of two standard document types, `User` and `UserCreate` (which are used in the default user management screens and backing beans) to add our new field:

```
<extension target="org.nuxeo.ecm.platform.types.TypeService" point="types">

  <type id="User" coretype="User">
    <label>User</label>
    <icon>/icons/user.gif</icon>
    <default-view>view_user</default-view>
    <layout>
      <widget schemaname="myuser" fieldname="username"
        jsfcomponent="h:inputTextReadOnly" />
      <widget schemaname="myuser" fieldname="firstName"
        jsfcomponent="h:inputText" />
      <widget schemaname="myuser" fieldname="lastName"
        jsfcomponent="h:inputText" />
      <widget schemaname="myuser" fieldname="email"
        jsfcomponent="h:inputText" />
      <widget schemaname="myuser" fieldname="company"
        jsfcomponent="h:inputText" />
      <widget schemaname="myuser" fieldname="petName"
        jsfcomponent="h:inputText" />
    </layout>
  </type>

  <type id="UserCreate" coretype="UserCreate">
    <label>UserCreate</label>
    <icon>/icons/user.gif</icon>
    <default-view>create_user</default-view>
    <layout>
      <widget schemaname="myuser" fieldname="username"
        jsfcomponent="h:inputText" required="true" />
      <widget schemaname="myuser" fieldname="password"
        jsfcomponent="h:inputSecret" required="true" />
      <widget schemaname="myuser" fieldname="firstName"
        jsfcomponent="h:inputText" />
      <widget schemaname="myuser" fieldname="lastName"
        jsfcomponent="h:inputText" />
      <widget schemaname="myuser" fieldname="email"
        jsfcomponent="h:inputText" required="true" />
      <widget schemaname="myuser" fieldname="company"
        jsfcomponent="h:inputText" />
      <widget schemaname="myuser" fieldname="petName"
        jsfcomponent="h:inputText" />
    </layout>
  </type>

</extension>
```

13.3. Authentication

13.3.1. Authentication Framework Overview

Nuxeo Authentication is based on the JAAS standard.

Authentication infrastructure is based on 2 main components :

- a JAAS Login Module: NuxeoLoginModule
- a Web Filter: NuxeoAuthenticationFilter

Users and groups and managed via the UserManagerService that handles the indirection to users and groups directories (SQL or LDAP).

Nuxeo authentication framework is pluggable so that you can contribute new plugin and don't have to rewrite and reconfigure a complete JAAS infrastructure.

13.3.2. Pluggable JAAS Login Module

NuxeoLoginModule is a JAAS LoginModule. It is responsible for handling all login call within Nuxeo's security domains:

- `nuxeo-ecm`: for the service stack and the core
- `nuxeo-ecm-web`: for the web application on the top of the service stack

On JBoss application server, the JBoss Client Login module is used to propagate security between the web part and the service stack.

Here is the default JBoss security configuration:

```
<domain name="nuxeo-ecm-web">
     <login-module code = "org.nuxeo.ecm.platform.login.NuxeoLoginModule"
            flag = "required">
             <option name="principalClassName">org.nuxeo.ecm.platform.login.NuxeoPrincipal</option>
             <option name="useUserIdentificationInfoCB">true</option>
         </login-module>
      <login-module code="org.jboss.security.ClientLoginModule" flag="required">
           <option name="password-stacking">true</option>
           <option name="restore-login-identity">true</option>
           <option name="multi-threaded">true</option>
       </login-module>
</domain>
<domain name="nuxeo-ecm">
       <login-module code = "org.nuxeo.ecm.platform.login.NuxeoLoginModule" flag = "required">
             <option name="principalClassName">org.nuxeo.ecm.platform.login.NuxeoPrincipal</option>
             <option name="useUserIdentificationInfoCB">true</option>
        </login-module>
</domain>
```

As shown by this configuration, the principals returned by NuxeoLoginModule is `org.nuxeo.ecm.platform.login.NuxeoPrincipal`.

Each protected service declares the `nuxeo-ecm` security domain

```
<?xml version="1.0" encoding="UTF-8"?>
<jboss>
  <enterprise-beans>
    <session>
        <ejb-name>DocumentManagerBean</ejb-name>
        <security-domain>nuxeo-ecm</security-domain>
    </session>
  </enterprise-beans>
</jboss>
```

13.3.2.1. NuxeoLoginModule

NuxeoLoginModule mainly handles 2 tasks:

- login user

  This means extract information from the CallBack stack and validate identity.

  NuxeoLoginModule supports several types of CallBacks (including Nuxeo specific CallBack) and uses a plugin system to be able to validate user identity in a pluggable way.

- Principal creation

  For that NuxeoLoginModule uses Nuxeo UserManager service that does the indirection to the users/groups directories.

When used in conjonction with UserIdentificationInfoCallback (Nuxeo custom CallBack system), the

LoginModule will choose the right LoginPlugin according to the CallBack information.

13.3.2.2. NuxeoLoginModule Plugins

Because validating User identity can be more complex that just checking login/password, NuxeoLoginModule exposes an extension point to contribute new LoginPlugins

Each LoginPlugin has to implement the `org.nuxeo.ecm.platform.login.LoginPlugin interface`.

Main method is:

```
String validatedUserIdentity(UserIdentificationInfo userIdent)
```

that is used to validate the UserIdentificationInfo.

Typically, default implementation will extract Login/Password from UserIdentificationInfo and call the checkUsernamePassword against the UserManager that will validate this information against the users directory.

Other plugins can use other informations carried by UserIdentificationInfo (token, ticket ...) to validate the identity against an external SSO system. The UserIdentificationInfo also carries the LoginModule plugin name that must be used to validate identity. Even if technically, a lot of SSO system could be implemented using this plugin system, most SSO implementations have be moved to the Authentication Plugin at the Web Filter level, because they need a http dialog.

For now, the NuxeoLoginModule has only two way to handle validateUserIdentity:

- default

  Uses UserManager

- Trusted_LM

  This plugin assumes the user identity has already been validated by the authentication filter, so the validatedUserIdentity will always return true. Using this LoginModule plugin, a user will be logged if the user exists in the UserManager. This plugin is used for most SSO system in conjonction with a Authentication plugin that will actually do the work of validating password or token.

13.3.2.3. Remote Login to the EJB layer

The Login system can be used via Remote EJB calls.

You can login as System user:

```
LoginContext lc = Framework.login();
// do some service calls
lc.logout();
```

You can login using user / password:

```
LoginContext lc = Framework.login(userName, password);
// do some service calls
lc.logout();
```

You can also call the login method and pass it directly a CallBackHandler. This can be used in conjonction with org.nuxeo.ecm.platform.api.login.UserIdentificationInfoCallbackHandler.

13.3.3. Pluggable Web Authentication Filter

The Web Authentication filter is responsible for:

- guarding access to web resources

  The filter can be parametrized to guard urls with a given pattern

- finding the right plugin to get user identification information

  This can be getting a userName/Password, getting a token in a cookie or a header, redirecting user to another authentication server.

- create the LoginContext

  This means creating the needed callBacks and call the JAAS Login

- store and reestablish login context

  In order to avoid recreating a login context for each request, the LoginContext is cached.

13.3.3.1. NuxeoAuthenticationFilter

The NuxeoAuthenticationFilter is one of the top level filter in Nuxeo Web Filters stack.

For each request it will try to find a existing LoginContext and create a RequestWrapper that will carry the NuxeoPrincipal.

If no existing LoginContext is found it will try to prompt the client for authentication information and will establish the login context.

If order to execute the task of prompting the client and retrieving UserIndetificationInfo, the filter will rely on a set of configured plugins.

Each plugin must:

- Implement org.nuxeo.ecm.platform.ui.web.auth.interfaces.NuxeoAuthenticationPlugin

  The two main methods are:

```
Boolean handleLoginPrompt(HttpServletRequest httpRequest,HttpServletResponse httpResponse, String baseURL);
UserIdentificationInfo handleRetrieveIdentity(HttpServletRequest httpRequest, HttpServletResponse httpResponse);
```

- Define the LoginModule plugin to use if needed

  Typically, SSO AuthenticationPlugin will do all the work and will use the Trusted_LM LoginModule Plugin.

- Define if stating URL must be saved

  AuthenticationPlugins that uses HTTP redirect in order to do the login prompt will let the Filter store the first accessed URL in order to cleanly redirect the user to the page he asked after the authentication is successful.

Additionnaly, AuthenticationPlugin can also implement the `org.nuxeo.ecm.platform.ui.web.auth.interfaces.NuxeoAuthenticationPluginLogoutExtension` interface if a specific processing must be done when logging out.

Here is a sample XML descriptor for registering an AuthenticationPlugin:

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.platform.ui.web.auth.defaultConfig">
  <extension
    target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
    point="authenticators">
```

```
    <authenticationPlugin name="FORM_AUTH" enabled="true"
      class="org.nuxeo.ecm.platform.ui.web.auth.plugins.FormAuthenticator">
      <needStartingURLSaving>true</needStartingURLSaving>
      <parameters>
        <parameter name="LoginPage">login.jsp</parameter>
        <parameter name="UsernameKey">user_name</parameter>
        <parameter name="PasswordKey">user_password</parameter>
      </parameters>
    </authenticationPlugin>
   </extension>
</component>
```

As you can see in the above example, the descriptor contains the parameters tag that can be used to embed arbitrary additional configuration that will be specific to a given AuthenticationPlugin. In the above example, it is used to define the field names and the JSP file used for form based authentication.

NuxeoAuthenticationFilter supports several authentication system. This is, for example, useful for having users using Form based authentication and having RSS clients using Basic Authentication. Because of that AuthenticationPlugin must be ordered. For that purpose, NuxeoAuthenticationFilter uses a dedicated extension point that let you define the AuthenticationChain.

```
<component name="Anonymous.auth.activation">
   <require>org.nuxeo.ecm.platform.ui.web.auth.defaultConfig</require>
    <extension
      target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
      point="chain">

    <authenticationChain>
      <plugins>
        <plugin>BASIC_AUTH</plugin>
        <plugin>ANONYMOUS_AUTH</plugin>
        <plugin>FORM_AUTH</plugin>
      </plugins>
    </authenticationChain>
  </extension>
</component>
```

The NuxeoAuthenticationFilter will use this chain to trigger the login prompt. When authentication is needed, the Filter will call, in a first round, the `handleRetrieveIdentity` method on all the plugins in the order of the authentication chain and then, in a second round, the `handleLoginPrompt` method in the same order on all the plugins if the authentication could not be achieved in the first round. The aim is to have as much automatic authentications as possible, that's why all the manual authentications (those which need a prompt) are done in a second round.

Some authentication plugins may choose to trigger or not the LoginPrompt depending on the situation. For example: the BasicAuthentication plugin generates the login prompt (in the case of the BasicAuthentication plugin the login prompt is an HTTP basic authentication which takes the form of a popup) only for specific URLs used for RSS feeds or restlet calls. This allows the platform to be easily called by Restlets and RSS clients without bothering browser clients who are presented with web forms to authenticate.

13.3.3.2. Built-in Authentication Plugins

NuxeoAuthenticationFilter comes with two built-in authentication plugins:

- FORM_AUTH: Form based Authentication

  This is a standard form based authentication. Current implementation let you configure the name of the Login and Password fields, and the name of the page used to display the login page

- BASIC_AUTH: Basic HTTP Authentication

  This plugin supports standard HTTP Basic Authentication. By default, this plugin only generates the authentication prompt on configured URLs.

There are also additional components that provides other Authentication plugins (see below).

13.3.3.3. Additional Authentication Plugins

Nuxeo provides a set of other authentication plugins that are not installed by default with the standard Nuxeo EP setup. These plugins can be downloaded and installed separately.

13.3.3.3.1. CAS2 Authentication

This plugin implements a client for CAS SSO system (Central Authentication System). It can be configured to use a CAS proxy. It has been tested and reported to work with CAS V2.

It's easy to test this plugin by installing the JA-SIG Central Authentication Service Open Source CAS server.

To install this authentication plugin, you need to:

- be sure that there is a CAS server already setup and running
- download the nuxeo-platform-login-cas2 plugin
- put it in $JBOSS_HOME/server/default/deploy/nuxeo.ear/plugins and restart the server
- configure the CAS2 descriptor
- put CAS2 plugin into the authentication chain

In order to configure CAS2 Auth, you need to create an XML configuration file into `$JBOSS_HOME/server/default/deploy/nuxeo.ear/config`

Here is a sample file named `CAS2-config.xml`.

```
<component name="MyAPP.Cas2SSO">
   <require>org.nuxeo.ecm.platform.ui.web.auth.defaultConfig</require>
   <require>org.nuxeo.ecm.platform.login.Cas2SSO</require>

   <!--  configure you CAS server parameters -->
   <extension
      target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
      point="authenticators">
      <authenticationPlugin
              name="CAS2_AUTH">
       <loginModulePlugin>Trusting_LM</loginModulePlugin>
       <parameters>
         <parameter name="ticketKey">ticket</parameter>
         <parameter name="appURL">http://127.0.0.1:8080/nuxeo/nxstartup.faces</parameter>
         <parameter name="serviceLoginURL">http://127.0.0.1:8080/cas/login</parameter>
         <parameter name="serviceValidateURL">http://127.0.0.1:8080/cas/serviceValidate</parameter>
         <parameter name="serviceKey">service</parameter>
         <parameter name="logoutURL">http://127.0.0.1:8080/cas/logout</parameter>
       </parameters>
      </authenticationPlugin>
  </extension>

   <!-- Include CAS2 into authentication chain -->
    <extension
      target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
      point="chain">
    <authenticationChain>
      <plugins>
        <plugin>BASIC_AUTH</plugin>
        <plugin>CAS2_AUTH</plugin>
      </plugins>
    </authenticationChain>
  </extension>
</component>
```

Tip: If while authenticating on the CAS server you get in the logs the following exception, it simply means that the user JOEUSER does not exist in the Nuxeo directory and does not mean that the CAS process is not working.

```
ERROR [org.nuxeo.ecm.platform.login.NuxeoLoginModule] createIdentity failed
javax.security.auth.login.LoginException: principal JOEUSER does not exist
        at org.nuxeo.ecm.platform.login.NuxeoLoginModule.createIdentity(NuxeoLoginModule.java:304)
        at org.nuxeo.ecm.platform.login.NuxeoLoginModule.validateUserIdentity(NuxeoLoginModule.java:362)
        at org.nuxeo.ecm.platform.login.NuxeoLoginModule.getPrincipal(NuxeoLoginModule.java:216)
        at org.nuxeo.ecm.platform.login.NuxeoLoginModule.login(NuxeoLoginModule.java:271)
        at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
        at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39)
        at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25)
        at java.lang.reflect.Method.invoke(Method.java:585)
```

```
        at javax.security.auth.login.LoginContext.invoke(LoginContext.java:769)
        at javax.security.auth.login.LoginContext.access$000(LoginContext.java:186)
        at javax.security.auth.login.LoginContext$4.run(LoginContext.java:683)
        at java.security.AccessController.doPrivileged(Native Method)
        at javax.security.auth.login.LoginContext.invokePriv(LoginContext.java:680)
        at javax.security.auth.login.LoginContext.login(LoginContext.java:579)
        at org.nuxeo.ecm.platform.ui.web.auth.NuxeoAuthenticationFilter.doAuthenticate(NuxeoAuthenticationFilter
```

13.3.3.3.2. PROXY_AUTH: Proxy based Authentication

This plugin assumes Nuxeo is behind a authenticating reverse proxy that transmit user identity using HTTP headers. This modules has be used on projects that uses a apache reverse proxy using client certificates to authenticate. SSO system (Central Authentication System V2):

To install this authentication plugin, you need to:

- download the nuxeo-platform-login-mod_sso plugin
- put it in `$JBOSS_HOME/server/default/deploy/nuxeo.ear/plugins` and restart the server
- configure the plugin via an XML descriptor
- put the plugin into the authentication chain

In order to configure this plugin, you need to create an XML configuration file into `$JBOSS_HOME/server/default/deploy/nuxeo.ear/config`

Here is a sample file named proxy-auth-config.xml.

```
<component name="MyAPP.Mod_sso">
   <require>org.nuxeo.ecm.platform.ui.web.auth.defaultConfig</require>
   <require>org.nuxeo.ecm.platform.login.Proxy</require>

   <extension
      target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
      point="authenticators">
      <authenticationPlugin
              name="PROXY_AUTH">
       <loginModulePlugin>Trusting_LM</loginModulePlugin>
       <parameters>
         <!-- configure here the name of the http header that is used to retrieve user identity -->
         <parameter name="ssoHeaderName">remote_user</parameter>
       </parameters>
      </authenticationPlugin>
  </extension>

   <!-- Include Proxy Auth into authentication chain -->
    <extension
      target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
      point="chain">
    <authenticationChain>
      <plugins>
        <!--  Keep basic Auth at top of Auth chain to support RSS access via BasicAuth -->
        <plugin>BASIC_AUTH</plugin>
        <plugin>PROXY_AUTH</plugin>
      </plugins>
    </authenticationChain>
  </extension>
</component>
```

13.3.3.3.3. NTLM_AUTH: NTLM and IE challenge/response authentication

This plugin uses JCIFS to handle NTLM authentication.

This plugging was partially contributed by Nuxeo EP users and has been reported to work by several users.

If you have troubles with latest version of IE on POST requests, please see JCIFS instructions on that (http://jcifs.samba.org/src/docs/ntlmhttpauth.html#post).

To install this authentication plugin, you need to :

- download the nuxeo-platform-login-ntlm plugin

- put it in `$JBOSS_HOME/server/default/deploy/nuxeo.ear/plugins` and restart the server
- configure the plugin via an XML descriptor
- put the plugin into the authentication chain

In order to configure this plugin, you need to create an XML configuration file into `$JBOSS_HOME/server/default/deploy/nuxeo.ear/config`

Here is a sample file named ntlm-auth-config.xml.

```
<component name="MyAPP.NTLM">
   <require>org.nuxeo.ecm.platform.ui.web.auth.defaultConfig</require>
   <require>org.nuxeo.ecm.platform.login.NTLM</require>

   <extension
      target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
      point="authenticators">
      <authenticationPlugin
              name="NTLM_AUTH">
       <loginModulePlugin>Trusting_LM</loginModulePlugin>
       <parameters>
          <!-- Add here parameters for you domain, please ee http://jcifs.samba.org/src/docs/ntlmhttpauth.html
         <parameter name="jcifs.http.domainController">MyControler</parameter>
          -->
       </parameters>

      </authenticationPlugin>
  </extension>

   <!-- Include NTLM Auth into authentication chain -->
    <extension
      target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
      point="chain">
    <authenticationChain>
      <plugins>
        <plugin>BASIC_AUTH</plugin>
        <plugin>NTLM_AUTH</plugin>
        <plugin>FORM_AUTH</plugin>
      </plugins>
    </authenticationChain>
  </extension>
</component>
```

13.3.3.3.4. PORTAL_AUTH: SSO implementation for portal clients

This plugin provides a way to handle identity propagation between an external application and Nuxeo. It was coded in order to propagate user identify between a JSR168 portal and a Nuxeo server. See the Nuxeo-Http-client-library for more information.

To install this authentication plugin, you need to :

- download the nuxeo-platform-login-portal-sso plugin
- put it in `$JBOSS_HOME/server/default/deploy/nuxeo.ear/plugins` and restart the server
- configure the plugin via an XML descriptor
- put the plugin into the authentication chain

In order to configure this plugin, you need to create an XML configuration file into `$JBOSS_HOME/server/default/deploy/nuxeo.ear/config`

Here is a sample file named `portal-auth-config.xml`.

```
<component name="MyAPP.postal_sso">
   <require>org.nuxeo.ecm.platform.ui.web.auth.defaultConfig</require>
   <require>org.nuxeo.ecm.platform.login.Portal</require>

   <extension
      target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
      point="authenticators">
      <authenticationPlugin
```

```
              name="PORTAL_AUTH">
       <loginModulePlugin>Trusting_LM</loginModulePlugin>
       <parameters>
         <!-- define here shared secret between the portal and Nuxeo server -->
         <parameter name="secret">nuxeo5secretkey</parameter>
         <parameter name="maxAge">3600</parameter>
       </parameters>
      </authenticationPlugin>
  </extension>

   <!-- Include Portal Auth into authentication chain -->
    <extension
      target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
      point="chain">
    <authenticationChain>
      <plugins>
        <!--  Keep basic Auth at top of Auth chain to support RSS access via BasicAuth -->
        <plugin>BASIC_AUTH</plugin>
        <plugin>PORTAL_AUTH</plugin>
        <plugin>FORM_AUTH</plugin>
      </plugins>
    </authenticationChain>
  </extension>
</component>
```

13.3.3.3.5. ANONYMOUS_AUTH: Anonymous authentication plugin

This plugin provides anonymous authentication. Users are automatically logged as a configurable Anonymous user. This modules also includes additional actions (to be able to login when already logged as Anonymous) and a dedicated Exception handling (to automatically redirect Anonymous users to login screen after a security error).

To install this authentication plugin, you need to :

- download the nuxeo-platform-login-anonymous plugin
- put it in `$JBOSS_HOME/server/default/deploy/nuxeo.ear/plugins` and restart the server
- configure the plugin via an XML descriptor (define who the anonymous user will be)
- put the plugin into the authentication chain

In order to configure this plugin, you need to create an XML configuration file into `$JBOSS_HOME/server/default/deploy/nuxeo.ear/config`

Here is a sample file named `anonymous-auth-config.xml`.

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.platform.login.anonymous.config">

  <!-- Make sure these components are read first -->
  <require>org.nuxeo.ecm.platform.ui.web.auth.defaultConfig</require>
  <require>org.nuxeo.ecm.platform.login.anonymous</require>

  <!-- Add an Anonymous user -->
  <extension target="org.nuxeo.ecm.platform.usermanager.UserService"
    point="userManager">
    <userManager>
      <users>
        <anonymousUser id="Guest">
          <property name="firstName">Guest</property>
          <property name="lastName">User</property>
        </anonymousUser>
      </users>
    </userManager>
  </extension>

  <!-- Override the default authentication chain present in
    nuxeo-platform-ui-web to add ANONYMOUS_AUTH. -->
  <extension
    target="org.nuxeo.ecm.platform.ui.web.auth.service.PluggableAuthenticationService"
    point="chain">
    <authenticationChain>
      <plugins>
        <plugin>BASIC_AUTH</plugin>
        <plugin>ANONYMOUS_AUTH</plugin>
        <plugin>FORM_AUTH</plugin>
      </plugins>
    </authenticationChain>
  </extension>
</component>
```

Chapter 14. Security Policy Service

14.1. Introduction

The Security Service provides an extension point to plug custom security policies that do not rely on the document security settings. For instance, it can be used to define permissions according to the document metadata, or to the logger user information.

14.2. Architecture

Policies are set in two places:

- in the core where custom checks of permission can be performed before resolving the document security using its ACP.
- in the search where the query can be patched with custom constraints before processing it.

Policies are checked in the order given as a parameter when registering them.

On the core side, they are used when checking permissions: they can grant or deny access, in case following policies - as well as the default security check relying on the ACP set on the document - will be ignored. They can also return an undefined access, in case following policy checks will continue.

When defining a custom policy for the Read permission, queries to the search service have to be adapted to have the same constraints: otherwise some queries will not return documents that the user can see, or will return documents that the user cannot see. Search policies are used when performing any search.

These policies are set on different services and follow different interfaces. When deploying on multi machines environment, where search and core may be hosted on different machines, the core and search policies will have to be deployed on the corresponding machine.

14.3. Policy contributions

14.3.1. Core policy contribution

To register a core security policy, you need to write a contribution specifying the class name of your implementation.

Example 14.1. Sample core security policy contribution

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.core.securityPolicy.accessLevelContrib">

  <extension target="org.nuxeo.ecm.core.security.SecurityService"
    point="policies">

    <policy name="accessLevel"
      class="org.nuxeo.ecm.core.security.AccessLevelSecurityPolicy" order="0" />

  </extension>

</component>
```

Here is a sample contribution: core policies have to follow the org.nuxeo.ecm.core.security.SecurityPolicy interface.

```
public class AccessLevelSecurityPolicy extends AbstractSecurityPolicy {

    public Access checkPermission(Document doc, ACP mergedAcp,
            Principal principal, String permission,
            String[] resolvedPermissions, String[] additionalPrincipals) {
        Access access = Access.UNKNOWN;
        try {
            if (principal instanceof NuxeoPrincipal) {
                DocumentModel userModel = ((NuxeoPrincipal) principal).getModel();
                if (userModel != null) {
                    Long accessLevel = (Long) userModel.getPropertyValue("user:accessLevel");
                    Long securityLevel = (Long) doc.getPropertyValue("sp:securityLevel");
                    if (accessLevel < securityLevel) {
                        access = Access.DENY;
                    }
                }
            }
        } catch (Exception e) {
        }
        return access;
    }
}
```

Here a specific check is done on all documents, regardless of the permission being checked. A property on the document (securityLevel) is checked against a property set on the user model (accessLevel).

Note that the "unknown" access is returned when the policy should not interfere with the standard security check.

14.3.2. Search policy contribution

To register a search policy contribution, you need to write a contribution specifying the class name of your implementation.

Example 14.2. Sample search policy contribution

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.platform.searchPolicy.accessLevelContrib">

  <extension target="org.nuxeo.ecm.core.search.service.SearchServiceImpl"
    point="policies">

    <policy name="access"
      class="org.nuxeo.ecm.core.search.security.AccessLevelSearchPolicy"
      order="20" />

  </extension>

</component>
```

Here is a sample contribution: search policies have to follow the org.nuxeo.ecm.core.search.api.security.SearchPolicy interface.

```
public class AccessLevelSearchPolicy extends AbstractSearchPolicy {

    private static final Log log = LogFactory.getLog(AccessLevelSearchPolicy.class);

    public ComposedNXQuery applyPolicy(ComposedNXQuery nxqlQuery)
            throws SearchException {
        SQLQuery query = nxqlQuery.getQuery();
        NuxeoPrincipal principal = null;
        SearchPrincipal sPrincipal = nxqlQuery.getSearchPrincipal();
        if (sPrincipal != null) {
            principal = (NuxeoPrincipal) sPrincipal.getOriginalPrincipal();
        }
        if (principal == null) {
            return nxqlQuery;
        }
        Predicate add = new Predicate(new Reference("sp:securityLevel"),
                Operator.LTEQ, new IntegerLiteral(
                        (Long) principal.getModel().getProperty("user",
                                "accessLevel")));
        if (!query.getWhereClause().toString().contains(add.toString())) {
            WhereClause wc;
            if (query.getWhereClause() != null) {
                wc = new WhereClause(new Predicate(
```

```
                        query.getWhereClause().predicate, Operator.AND, add));
            } else {
                wc = new WhereClause(add);
            }
            query = new SQLQuery(query.getSelectClause(),
                    query.getFromClause(), wc, query.groupBy, query.having,
                    query.getOrderByClause());
        }
        return new ComposedNXQueryImpl(query, sPrincipal);
    }
}
```

As a check is done on the core side regardless of the permission, it will also apply for the "Read" permission which is checked when navigating to a document. Without this policy, searches would return documents that the user may not be able to see, and woul generate a security exception when the user clicks on the document link.

The same check than performed on the core is added here as a contraint to the search query. Note that the principal attached to the query is kept: it is used to perform standard security checks, as the search indexes the document ACP.

For additional information about NXQL queries, please refer to the Section 11.3, “Programmatic Searching” section.

Chapter 15. Workflow & jBPM

15.1. Workflow in Nuxeo 5.1

15.1.1. Deploying process definitions

A workflow definition is designed for a given workflow engine backend and not for the Nuxeo workflow service itself. Nuxeo workflow doesn't specify a new process definition language. Thus it has no Nuxeo specificities speaking of format or process definition language. For instance, if you use jBPM as a backend with Nuxeo 5 then the workflow definition should be a standard jpdl file that you may have designed using your favorite editor or still if you are using Shark as a backend then the workflow definition will be a standard WFMC process definition.

Once your workflow definition has been designed and is ready you can deploy it in Nuxeo workflow. Of course, the target workflow engine backend plugin should be deployed and registered against the workflow service.

15.1.1.1. Using extension points

The Nuxeo workflow service provides a dedicated extension point for workflow definition deployment. The extension point is called **definition.**

In this case, the workflow definition will be deployed at application server deployment time (For now, this is the case when the application server is starting up since hot deployment is not yet possible using Nuxeo Runtime at the time of writing this document).It means this way of deploying workflow definition is not suitable for all cases. See the next subsections for other ways of deploying workflow definitions.

Below is an example of a jPDL workflow definition contribution for the jBPM backend. This XML fragment would be defined in a contribution registered as a component in a bundle:

```
<?xml version="1.0"?>

<component name="com.company.workflow.sample.contributions">

  <extension target="org.nuxeo.ecm.platform.workflow.service.WorkflowService"
             point="definition">
    <definition>
      <engineName>jbpm</engineName>
      <mimetype>text/xml</mimetype>
      <definitionPath>workflows/process_definition.xml</definitionPath>
    </definition>
  </extension>

</component>
```

- **engineName:** name specified for the target backend at workflow service registration time (see workflow service backend extension point)
- **mimetype:** mimetype of the workflow definition. This is especially interesting in case of the format is binary. (serialization issue at deployment time)
- **definitionPath:** bundle relative path of the workklow definition to deploy.

In this situation here is how would look the tree:

```
com.company.workflow /
  META-INF /
    workflows /
      process_definition.xml
      MANIFEST.MF
  OSGI-INF /
    workflow-definitions-contrib.xml
```

15.2. Workflow from Nuxeo 5.2: the jBPM Service

15.2.1. Introduction

Nuxeo jBPM Service is a task and work flow management service based on jBPM. One can use this service without knowledge of jBPM, but for any advance use a good knowledge is necessary, please refer to jBPM documentation .

The service provides:

- an easy way to deploy process definitions (see Section 15.2.2.4, “deployer extension point”)
- 2 abstract classes to test process definition and the jBPM service using those definitions (see Section 15.2.5.4, “How to test a process ?”)
- A set of convention to facilitate manipulation of documents (see Section 15.2.3, “Document management”)
- Helpers class to create ActionHandler, AssignmentHandler, DecisionHandler and TaskControllerHandler that interact with repository, user manager and documents. (see Section 15.2.3, “Document management”)
- A set of hibernate queries that optimize queries on documents. (see Section 15.2.5.1, “How to optimize jBPM service ?”)
- 2 default workflows for document validation. (see Section 15.2.4, “Default processes”)

15.2.2. jBPM service configuration

15.2.2.1. Event and Notification

jBPM can be used for different use cases. The default Nuxeo contributions uses it as a simple workflow management on a single document. You could use it for complex BPM, SOA orchestration or pageflow management. All those different uses make it impossible to send meaningful event at the service level.

So we let the responsability of sending event to the user of the jBPM service. *It is the responsability of the user of the service to send event* .

For example, the default workflow will send a workflowStart event when the user presses the 'Start Workflow' button. (However, the jBPM process has already started and 2 tasks have already been completed at this stage). The default workflow also sends event (workflowTaskAssign) when a task is assigned to a pooled actors (This is not considered a task assignment for jBPM) and send a mail to each pooled actor.

If you use jBPM for workflow management on documents, we recommend that you use the event named in JbpmEventNames. You can check the default workflow, the TaskNotificationHandler or the JbpmActionsBean for code that sends event.

15.2.2.2. configurationPath and activeConfiguration extension points

Nuxeo configures jBPM in 2 steps . First, it creates named configurations using configurationPath extension point, such as:

```
<extension
    target="org.nuxeo.ecm.platform.jbpm.core.JbpmService"
    point="configurationPath">
  <configuration name="prod" path="config/jbpm.cfg.xml" />
</extension>
```

Finally, it gives the configuration to use:

```
<component
    name="org.nuxeo.ecm.platform.jbpm.core.ActiveConfiguration">
  <extension target="org.nuxeo.ecm.platform.jbpm.core.JbpmService"
    point="activeConfiguration">
    <configuration name="prod" />
  </extension>
</component>
```

If you want to use a different configuration you just need to overwrite the org.nuxeo.ecm.platform.jbpm.core.ActiveConfiguration component.

The default configuration uses a transactional persistent service. It also adds to the default hibernate configuration a nuxeo.hibernate.queries.hbm.xml file with nuxeo's specific queries.

Production settings

The default hibernate configuration, referenced in jbpm.cfg.xml, is not for production. For production, you should comment out the 'hibernate.hbm2ddl' property. You also should ask your DBA to finalize your database schema for your target database (Derby should not be used in production).

15.2.2.3. processDefinition extension point

To contribute a new process definition, contribute to the processDefinition extension point by giving the definition path relative to the root of the package and the name of the deployer.

The default contribution is:

```
<extension target="org.nuxeo.ecm.platform.jbpm.core.JbpmService"
    point="processDefinition">
  <processDefinition path="process/parallel-review.xml" deployer="ifChanged" />
  <processDefinition path="process/validation-review.xml" deployer="ifChanged" />
</extension>
```

15.2.2.4. deployer extension point

jBPM needs to keep all process definition versions in persistent storage, so a process may finish with the same definition it started with. When Nuxeo starts, the jBPM service needs to know if a process definition is to be deployed as a new version. The Deployer is responsible for process definition deployment, you choose how processes are deployed by choosing a Deployer.

The service comes with 3 Deployers:

ifChanged
This deployer deploys the process definition only if it wasn't deployed. To check it, it removes comments and whitespace from the process definition and compute the md5. If the previous definition had a different md5, then it is deployed.

never
This deployer never deploys a process definition.

always
This deployer always deploys a process definition.

You can contribute to new Deployer by contributing to the deployer extension-point. The contributed class needs to implement the ProcessDefinitionDeployer interface. It can also extend the AbstractProcessDefinitionDeployer that provides some helpers method.

The default contribution is:

```
<extension target="org.nuxeo.ecm.platform.jbpm.core.JbpmService"
```

```
    point="deployer">
  <deployer name="always"
    class="org.nuxeo.ecm.platform.jbpm.core.deployer.AlwaysDeployer" />
  <deployer name="never"
    class="org.nuxeo.ecm.platform.jbpm.core.deployer.NeverDeployer" />
  <deployer name="ifChanged"
    class="org.nuxeo.ecm.platform.jbpm.core.deployer.IfChangedDeployer" />
</extension>
```

15.2.2.5. securityPolicy extension point

The security extension point allows to contribute a security policy to a process definition. This security policy is used in the document related methods of the JbpmService.

The default contribution is:

```
<extension target="org.nuxeo.ecm.platform.jbpm.core.JbpmService"
  point="securityPolicy">

  <securityPolicy
    class="org.nuxeo.ecm.platform.jbpm.core.service.DefaultJbpmSecurityPolicy"
    for="review_parallel" />
  <securityPolicy
    class="org.nuxeo.ecm.platform.jbpm.core.service.DefaultJbpmSecurityPolicy"
    for="review_approbation" />

</extension>
```

This class follows the JbpmSecurityPolicy interface, and is used to define who is allowed to edit the process (administrator or initiator). Other permissions (who is allowed to add some reviewers for instance) is left out to the interface.

15.2.2.6. typeFilter extension point

The typeFilter extension point allows to bind process definition to a type of document. This is used in the default 'Workflow' tab to show what choices to offer to the user.

The default contribution is:

```
<extension target="org.nuxeo.ecm.platform.jbpm.core.JbpmService"
  point="typeFilter">
  <type name="Note">
    <processDefinition>review_parallel</processDefinition>
    <processDefinition>review_approbation</processDefinition>
  </type>
  <type name="File">
    <processDefinition>review_parallel</processDefinition>
    <processDefinition>review_approbation</processDefinition>
  </type>
</extension>
```

15.2.3. Document management

The jBPM service provides a set of methods to interact with a single document. We use a simple convention to attach a document to a process instance. 2 variables are added to the process instance:

documentId
The id of the document.

documentRepositoryName
The name of the document's repository.

The `VariableName` enum on the JbpmService names the different variables.

The abstract class `AbstractJbpmHandlerHelper` implements ActionHandler, AssignmentHandler, DecisionHandler and TaskControllerHandler. It provides useful methods when dealing with document and repository.

15.2.4. Default processes

15.2.4.1. The parallel review workflow

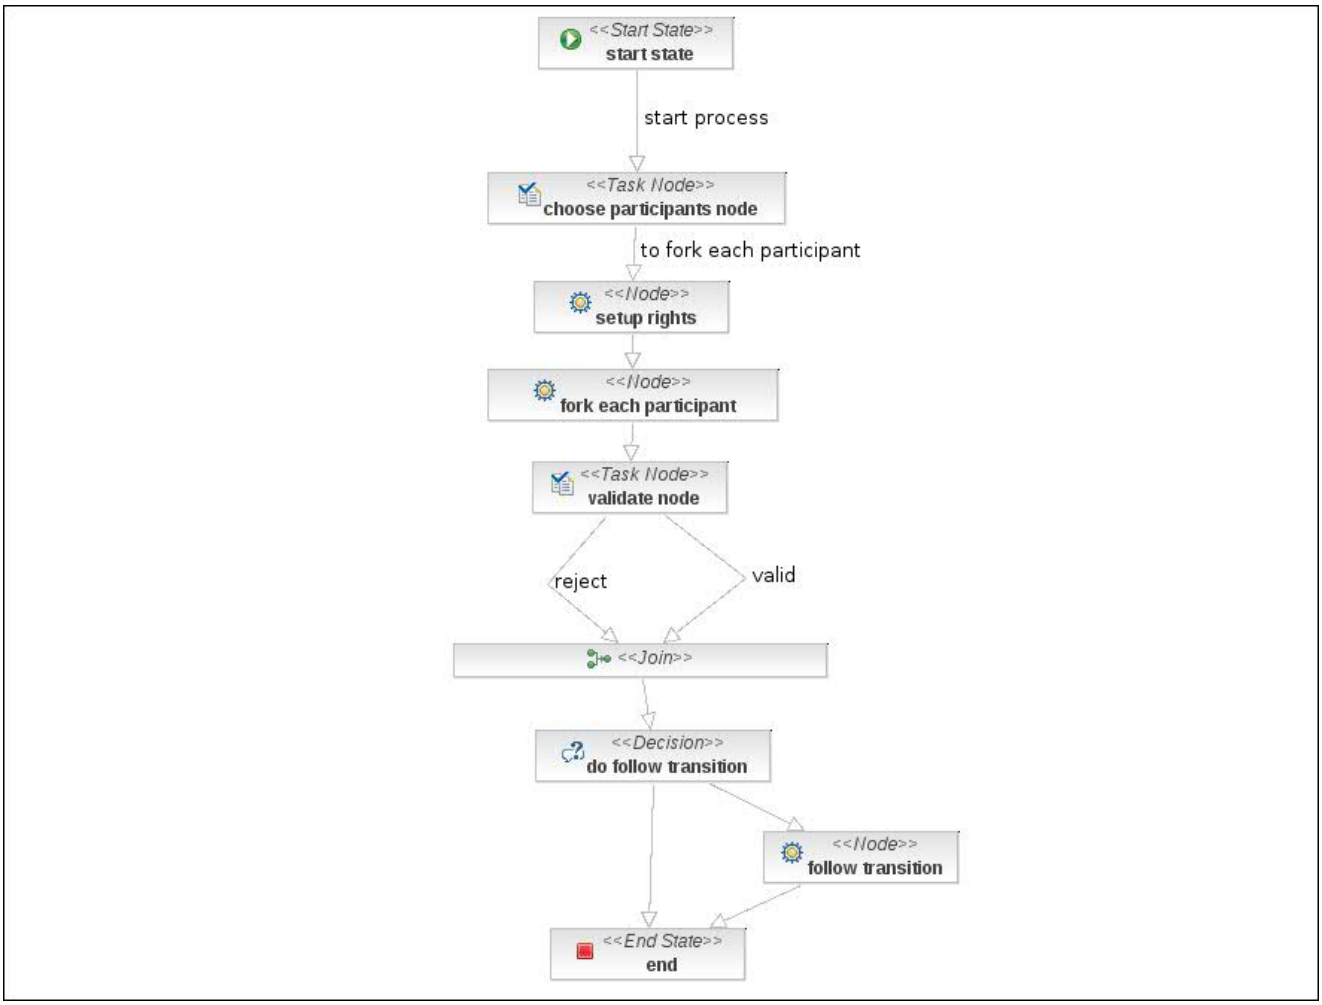

Description

Start state
When the process is started, 3 variables are added to the process: the document id and the repository name variable, the transition to follow(`endLifecycleTransition` variable). A start state task is also created to capture the initiator actor id. It is retrievable afterward as the 'initiator' swimlane.

Choose participants node
A task is created for the initiator. At the end of this task a 'participants' variable is added to the process. It is a list with one item per participant.

Set up rights
This is an action node. The process sets up the rights on the document so all participants can read it.

for each participant node
It uses an action node defined in the Nuxeo jBPM service. It creates a child token for each item of the list given as parameter in the process definition. In this case, a child token is created for each participant. On each child token, a variable named participant will hold the participant.

validate node
A task for the participant. If it is rejected, adds a variable named 'rejected' to the process. (The process is therefore considered rejected if only one participant rejects it)

join

A normal jBPM join. Waits for all participants to finish their tasks before resuming.

do follow transition
A decision node that checks if there is a 'rejected' variable. If there is, go to end. Otherwise go to follow transition node.

follow transition
The attached document follows the named transition

15.2.4.2. The validation review workflow

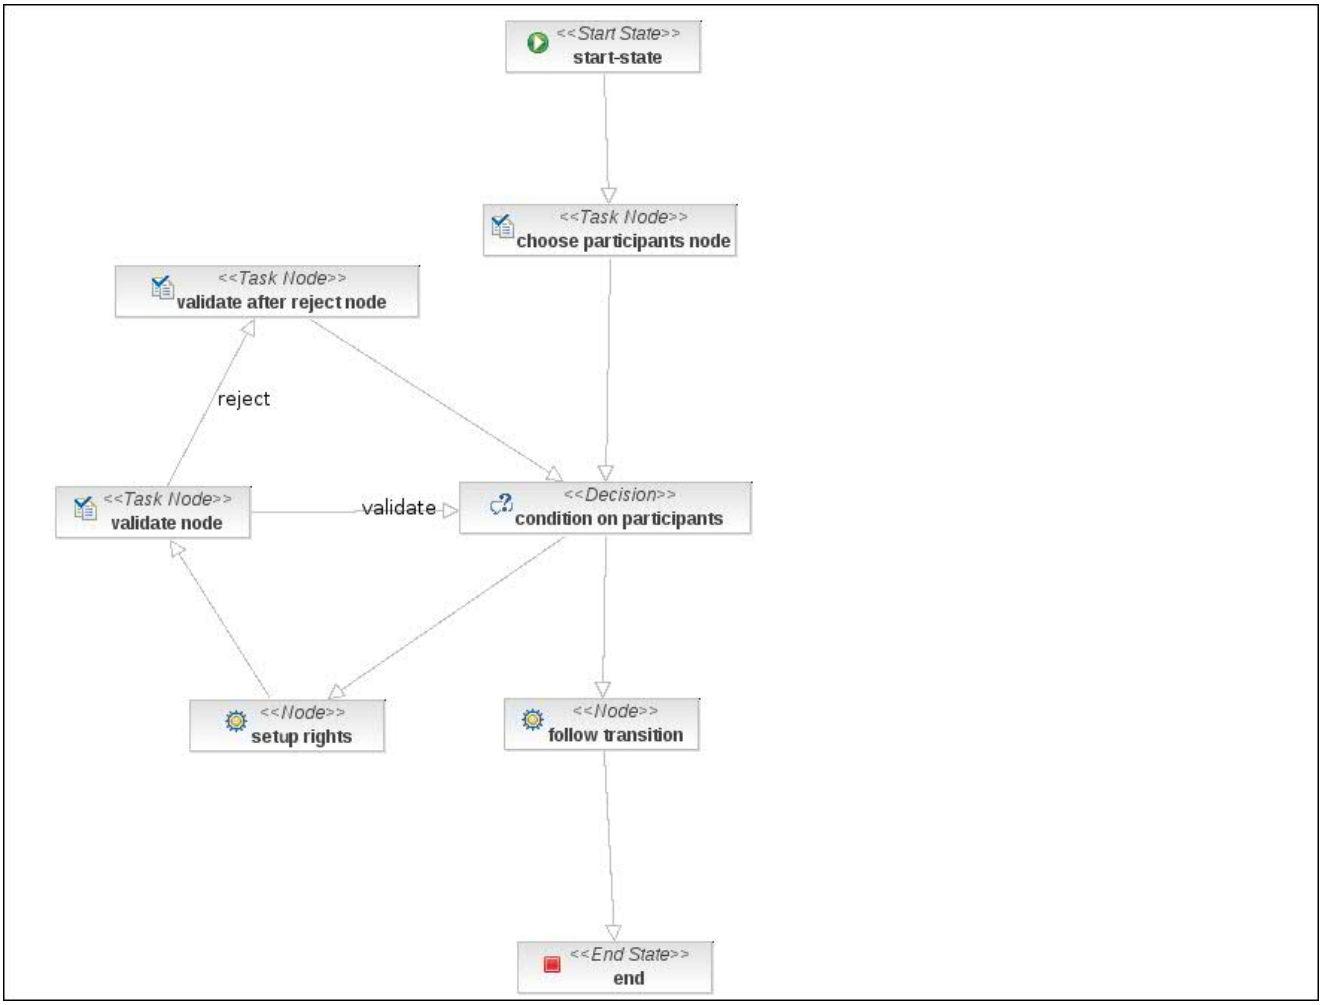

Description

start state
Same as for the parallel review process.

choose participants
Same as for the parallel review process.

Condition on participants
Check if there is an item in the participants list. If there isn't, then go to follow transition. Otherwise go to set up rights. Remove the first item of 'participants' list and put it in the participant variable. In any case, remember the previous participant in the previousActorId variable (the initiator if no previous participant).

Setup rights
Makes the document readable to the participants

validate node

Create a task for participant. If she validates, then go to condition on participants otherwise to validate after reject node.

validate after reject node
Create a task for previousActorId user. She has to validate the node.

follow transition
The document follows the transition set in the transition variable.

15.2.5. Nuxeo jBPM How-to

15.2.5.1. How to optimize jBPM service ?

The jBPM service methods on documents use the jBPM API. To find task related to a document it iterates over all the tasks and checks its variables.

If you have a lot of tasks, such a behavior can be too costly. We recommend in such a case to create a hibernate query and to call this query from your own jbpmOperation object. The module nuxeo-platform-jbpm-core comes with a `nuxeo.hibernate.queries.hbm.xml` that you can use. If you create your own query file, you need to add it in the hibernate configuration file.

If you need to query on property of the attached document, you can optimize the query by directly using the repository (this is valid only when using the sql backend). You need to create a class to map to the schema you want to query and add the hibernate annotation, you'll then be able to create a hbm query using the schema. You need to keep in mind that hibernate or the repository could be caching data, using such queries should be used carefully.

15.2.5.2. How to use jBPM as a task management service ?

The jBPM service can also be used as a simple task management service. The publication service uses it that way. It creates a task when someone needs to validates a publishing request. There is no process instance attached to it (but still a document attached to the task).

The code to create a task is:

```
TaskInstance ti = new TaskInstance();
ti.setPooledActors(new String[]{"some", "actors"});
Map<String, Serializable> variables = new HashMap<String, Serializable>();
variables.put(JbpmService.VariableName.documentId.name(), document.getId());
variables.put(JbpmService.VariableName.documentRepositoryName.name(), document.getRepositoryName());
ti.setVariables(variables);
ti.setName("my task");
ti.setCreate(new Date());
jbpmService().saveTaskInstances(Collections.singletonList(ti));
```

15.2.5.3. How to clean the database ?

jBPM keeps all tasks, processes and definitions in the database. If you don't need to keep an history of the processes, you have two ways to delete unwanted data.

You can use the event service to call the deleteProcess on all finished process periodically (Note: the default process in Nuxeo finishes only when all tasks are finished).

You can run a SQL procedure at the database level to remove the ended process and attached tasks.

15.2.5.4. How to test a process ?

Testing a process is done in two steps. The first step is to test the process definition itself. The second is to test the use of the process definition by the jbpm service and your other methods.

To test a process definition, you can extend the AbstractProcessDefinitionTest and implement the getProcessDefinitionResource() methods. It should return the path of the process definition. You can then use

the jbpmContext variable to interact with the process definition. You might have trouble if one of your action requests some of Nuxeo services. If your handlers extend the AbstractJbpmHandlerHelper, you can surround the logic inside a if(nuxeoHasStarted()){....}. This way no nuxeo services will be called inside unit test without services.

If you want to test a process definition and its interaction with nuxeo service, you can have a look at JbpmServiceTest. It creates a repository, deploys the user manager service and the sql directory.

Chapter 16. Document Versioning

16.1. Setting the version of a document

The versioning information are stored in the DocumentModel under the fields major_version and minor_version of the `uid.xsd` schema. Set the version as you would any other field:

```
documentModel.setProperty("uid","major_version",1);
documentModel.setProperty("uid","minor_version",1);
```

The field used for version is adaptable via the `properties` extension point of the `org.nuxeo.ecm.platform.versioning.service.VersioningService` component. It allows to define which properties should be used to set versions for a given document type.

```
<versioningProperties>
  <majorVersion>my:major_version</majorVersion>
  <minorVersion>my:minor_version</minorVersion>
  <documentType>File</documentType>
  <documentType>Note</documentType>
</versioningProperties>
```

16.2. Modifying automatically the version of a document

Event such as saving a document or following a life cycle transition can change the document's version number. To use this feature you need first to define which event should be listened to and then how the version number should behave.

The `CoreEventListenerService` is used to define which event to listen to. The default declaration is in `nuxeo-platform-versioning-core`:

```
<listener name="versioninglistener"
    class="org.nuxeo.ecm.platform.versioning.listeners.DocVersioningListener">
  <eventId>lifecycle_transition_event</eventId>
  <eventId>documentCreated</eventId>
  <eventId>beforeDocumentModification</eventId>
  <eventId>documentUpdated</eventId>
  <eventId>documentRestored</eventId>
</listener>
<listener name="versioningChangelistener"
  class="org.nuxeo.ecm.platform.versioning.listeners.VersioningChangeListener" />
```

The class `org.nuxeo.ecm.platform.versioning.listeners.DocVersioningListener` implements the behavior of the versioning when those events happen.

To modify the version of a document when a life cycle state is reach you need to define rules with the behavior(increment major, minor or do nothing) using the extension `rules` from `org.nuxeo.ecm.platform.versioning.service.VersioningService`:

```
<versioningRuleEdit name="sampleEditRuleProject" action="ask_user"
    lifecycleState="project">
  <option value="no_inc" default="true" />
  <option value="inc_minor" />
  <option value="inc_major" />
</versioningRuleEdit>
```

The default behavior for all type but File and Note is to increase the minor version for each life cycle change. You need to override one of the default rules if you add a new type. The order is important, the list of rules is read and the first match is used.

```
<versioningRuleEdit name="sampleEditRuleAnyState" action="ask_user"
      lifecycleState="*">
  <includeDocType>File</includeDocType>
  <includeDocType>Note</includeDocType>
  <includeDocType>MyNewType</includeDocType>
  <option value="no_inc" default="true" />
```

```
  <option value="inc_minor" />
  <option value="inc_major" />
</versioningRuleEdit>
```

16.3. Accessing document from previous version

The `CoreSession` or the seam component `documentVersioning` and `versionedActions` allow to access document from previous version. On `CoreSession`, the method `getVersions` return all the versions of a document.

16.4. The versioning service implementation

The versioning service manages the version inside Nuxeo. Two implementations are available:

- The custom Nuxeo service, more flexible and used by default
- The default JackRabbit implementation.

To modify this setting edit the file `config/default-versioning-config.xml`:

```
<!--property name="versioningService" value="org.nuxeo.ecm.core.repository.jcr.versioning.JCRVersioningService"/
<property name="versioningService" value="org.nuxeo.ecm.core.versioning.custom.CustomVersioningService"/>
```

This section shows how to use the versioning module, how to modify the version of a document automatically or manually and how to use a document from a previous version.

Chapter 17. Audit Service

17.1. Introduction

Audit service is used for logging and retrieving audit data into a datastore. Audit data are mainly coming from events.

17.2. Features

The audit service is logging creation/deletion/modification events. It is also possible to configure the service to log other events. For example, there is an addon, called nuxeo-platform-audit-web-access, that log web access.

17.3. Architecture

Audit service is mainly a datastore service. It defines a data record structure that will be used for storing audit information. The datastore is built over a relational database backend. The data record structure is defined in Java by the LogEntry and ExtendedInfo java classes. They are mapped onto the datastore using JPA (Java Persistence API) annotations. Audit service receive events from the Event service. Then the Audit service is filtering and converting them into log entries. The LogEntry class is mainly obtained from the DocumentMessage event type. Audit entries may also contain extended informations. These informations are extracted from the event message using EL (Expression Language) expression and stored into a map.

Note

Extended information map is a feature that is available since the 5.2 release. Prior releases was achieving extension by introducing specialized LogEntry types and OR mappings.

17.4. Retrieving entries

The following java snipset shows you how to retrieve entries associated with a document

```
..
  DocumentModel document = ...
  NXAuditEventService audit = Framework.getService(NXAuditEventService.class);
  List<LogEntry> entries = audit.getLogEntriesFor(document.getDocUUID());
..
```

You can also select entries using HQL language. The following snipset shows you how to retrieve entries for a whole document hierachy having a dublincore title.

Note

an extended information should be contributed to the audit service, extracting the dublincore title property from the document and storing it in the extended information map using the 'title' key.

```
..
  NXAuditEventService audit = Framework.getService(NXAuditEventService.class);
  List<LogEntry> entries = audit.nativeQueryLogs(
    "log.docPath like '/somefolder/%' and" +
    "log.extendedInfos['title'] is not null", 1, 10);
..
```

You may need to add some information to the audit datastore. Sending a core event is not the only way. You can invoke directly the audit service. The following java code snipset shows you how to do that.

```
..
  NXAuditEventService audit = Framework.getService(NXAuditEventService.class);
  LogEntry entry = new LogEntry();
  ..
  entry.setXXX(...);
  ..
  audit.addLogEntry(entry);
..
```

17.5. Contributing the audit service

17.5.1. Recording new events types

Logging other event types can be done by using an `event` extension point. Here is an example of how to define this extension point.

```
<extension point="event"
           target="org.nuxeo.ecm.platform.audit.service.NXAuditEventsService">
    <event name="documentCreated" />
    <event name="documentCreatedByCopy" />
    <event name="documentDuplicated" />
    <event name="documentMoved" />
    <event name="documentRemoved" />
    <event name="documentModified" />
    <event name="documentLocked" />
    <event name="documentUnlocked" />
    <event name="documentPublished" />
    <event name="documentSecurityUpdated" />
    <event name="documentUnPublished" />
    <event name="documentSubmitedForPublication" />
    <event name="documentPublicationRejected" />
    <event name="documentPublicationApproved" />
    <event name="lifecycle_transition_event" />
</extension>
```

17.5.2. Recording additional informations

Just after converting received DocumentMessage instance into the corresponding LogEntry instance, Audit service allows you to extract information from the handling context and to store them. To do this, you have to define an EL expression and associate it with a key. You can access to the following variables :

`message`
Document message describing the event

`source`
Document from which the event is from

`principal`
Identity of the event owner

The following XML snipset is an example of how to extract properties from the document model and store them into the extended information map.

```
<extension point="extendedInfo"
           target="org.nuxeo.ecm.platform.audit.service.NXAuditEventsService">
    <extendedInfo expression="${source.dublincore.title}" key="title" />
    <extendedInfo expression="${message.cacheKey}" key="key" />
    <extendedInfo expression="${principal.name}" key="user" />
</extension>
```

Chapter 18. Directories and Vocabularies

18.1. Introduction

TODO OG - General overview of the directory concept and goals

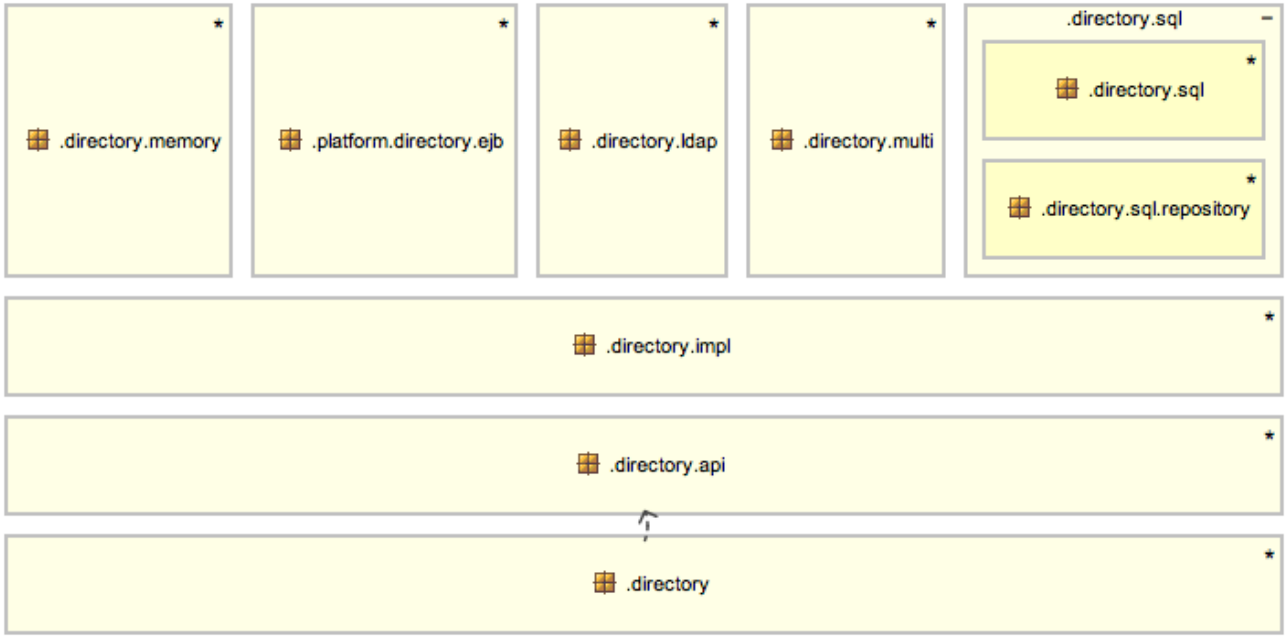

18.2. Directory with a Relational Database Management System (SQL) server as backend

The SQL server as storage backend is provided by org.nuxeo.ecm.directory.sql.* component.

```
<component name="org.nuxeo.ecm.directory.sql.storage">
  <implementation class="org.nuxeo.ecm.directory.sql.SQLDirectoryDescriptor" />

  <require>org.nuxeo.ecm.directory.sql.SQLDirectoryFactory</require>

  <extension target="org.nuxeo.ecm.directory.sql.SQLDirectoryFactory"
    point="directories">
    <directory name="userDirectory">
      <schema>user</schema>
      <dataSource>java:/nxsqldirectory</dataSource>
      <table>users</table>
      <idField>username</idField>
      <passwordField>password</passwordField>
      <autoincrementIdField>false</autoincrementIdField>
      <dataFile>users.csv</dataFile>
      <createTablePolicy>on_missing_columns</createTablePolicy>
      <querySizeLimit>15</querySizeLimit>
    </directory>
  </extension>
</component>
```

This code declares a **directories** node which defines a directory of users or of groups. The following information are given to describe the directory:

- `name`: name of the server which will be used in the declaration of the directories
- `schema`: namee of the schema describing the user attributes in the directory
- `dataSource`: type of storage for the directory. In this example, the HSQLDB is used. Other RDBMS like PostgreSQL can be used to store the datas by changing the local datasource.

- `table`: name of the table in which the data will be stored
- `idField`: the id field designs the primary key in the table, used for retrieving entries by id
- `password`: field from the table which contain the passwords, relative to the identfiant
- autoincrementIdField: boolean value which tells if the idField is automatically incremented - this value is most of the time at **false**, because the identifiant is a string.
- `dataFile`: file from which data are getting to populate the table. Be careful to follow the structure of the schema given above.
- `createTablePolicy`: indicates how the dataFile will be used to populate the table. Three values are allowed:
  - *never*: the dataFile is never used
  - *on_missing_columns*: the dataFile is used to create missing columns, it means at creation of the table or each time a new column is added, to follow the schema for example. Colums cannot be deleted.
  - *always*: the dataFile is used to create the table as each restart of the application server
- `querySizeLimit`: the maximum number of results that the queries on this directory should return; if there are more results than this, an exception will be raised

18.3. Directory with an LDAP server as backend

The LDAP server as storage backend is provided by the org.nuxeo.ecm.directory.ldap.* component

18.3.1. Server definition

First of all, LDAP servers have to be defined by adding a contribution to the **servers** extension point. The syntax is:

```
<extension target="org.nuxeo.ecm.directory.ldap.LDAPDirectoryFactory"
  point="servers">

  <server name="default">

    <ldapUrl>ldap://localhost:389</ldapUrl>
    <bindDn>cn=nuxeo5,ou=applications,dc=example,dc=com</bindDn>
    <bindPassword>changeme</bindPassword>
  </server>

</extension>
```

These information need to be provided to use an LDAP connection:

- `name`:name of the server which will be used in the declaration of the directories
- `ldapUrl`:address of the LDAP server. A single server declaration can point to a cluster of replicated servers. To leverage such a cluster and improve availibility, please provide one <ldapUrl/> tag for each replica of the cluster. ldaps is the convention to use TLS/SSL connection.
- `bindDn`: the Distinguished Name used to bind to the LDAP server
- `bindPassword`: corresponding password relative to the Distinguished Name for binding

  These credentials are used by Nuxeo5 to browse directory and create/modify entries.

18.3.2. Directory declaration

Once you have declared the server, you can define new LDAP directories, using the following syntax for its declaration:

```
<extension target="org.nuxeo.ecm.directory.ldap.LDAPDirectoryFactory"
  point="directories">

  <directory name="userDirectory">
    <server>default</server>
    <schema>user</schema>
    <idField>username</idField>
    <passwordField>password</passwordField>

    <searchBaseDn>ou=people,dc=example,dc=com</searchBaseDn>
    <searchClass>person</searchClass>
    <searchFilter>(&amp;(sn=toto*)(myCustomAttribute=somevalue))</searchFilter>
    <searchScope>onelevel</searchScope>
    <substringMatchType>subany</substringMatchType>

    <readOnly>false</readOnly>

    <cacheTimeout>3600</cacheTimeout>
    <cacheMaxSize>1000</cacheMaxSize>

    <creationBaseDn>ou=people,dc=example,dc=com</creationBaseDn>
    <creationClass>top</creationClass>
    <creationClass>person</creationClass>
    <creationClass>organizationalPerson</creationClass>
    <creationClass>inetOrgPerson</creationClass>

  </directory>

</extension>
```

The attributes are:

- **name**, *schema*, *idField* and *passwordField* are the same as for SQL directories
- `searchBaseDn`: entry point into the server's LDAP tree structure. Searches are only made below this root node
- `searchClass`: restricts the type of entries to return as result
- `searchFilter` :additional filter to restrict the search results
- `searchScope`: the scope of the search. It can take two values:
  - *onelevel*:search only under the current node.
  - *subtree*: search in the whole subtree. Use this parameter when the [people] branch is nested.
- `substringMatchType`: defines who the query is built using wildcard characters. Three different values can be provided:
  - *subany*: wildcards are added around the string to match (as *foo*)
  - *subinitial*: wildcard is added before the string (*bar)
  - *subfinal*: wildcard is added after the string (baz*). This is the default behaviour.
- `readOnly`: boolean value. This parameter allows to create new entries or modify existing ones in the LDAP server
- `cacheTimeout`: cache timeout in seconds
- `cacheMaxSize`: maximum number of cached entries before global invalidation

  To disable the cache, comment <cache* /> tags
- `creationBaseDn`: entry point in the server's LDAP tree structure where new entries will be created. This is useless to provided if readOnly attribute is set to false.

- `creationClass`: use as many tag as needed to specify which class are used to defined new people entries in LDAP server.

18.4. Handling references between directory entries

Directory references leverage two common ways of string relationship in LDAP directories

18.4.1. References defined by a many-to-many SQL table

TODO OG

18.4.2. Static reference as a dn-valued LDAP attribute

The static reference strategy is to apply when a multi-valued attribute stores the exhaustive list of distinguished names of reference entries, for example the uniqueMember of the groupOfUniqueNames object.

```
        <ldapReference field="members" directory="userDirectory"
          staticAttributeId="uniqueMember" />
```

The staticAttributeId attribute contains directly the value which can be read and manipulated.

18.4.3. Dynamic reference as a ldapUrl-valued LDAP attribute

The dynamic attribute strategy is used when a potentially multi-value attribute stores a LDAP URL intensively, for example the memberURL's attribute of the groupOfURL object class.

```
        <ldapReference field="members" directory="userDirectory"
          forceDnConsistencyCheck="false"
          dynamicAttributeId="memberURL" />
```

The value contained in `dynamidAttributeId` looks like ldap:///ou=groups,dc=example,dc=com??subtree?(cn=sub*) and will be resolved by dynamical queries to get all values. The `forceDnConsistencyCheck` attribute will check that the value got through the dynamic resolution correspond to the attended format. otherwise, the value will be ignored. Use this check when you are not sure of the validity of the distinguished name

18.4.4. Defining inverse references

Inverse references are defined with the following declarations:

```
      <references>
        <inverseReference field="groups" directory="groupDirectory"
          dualReferenceField="members" />
      </references>
```

This syntax should be understood as "the *member groups* value is an inverse reference on *groupDirectory* directory using *members* reference". It is the group directory that stores all members for a given group. So the groups of a member are retrieved by querying in which groups a member belongs to.

18.5. Combining multiple directories into a single virtual directory

Multi directories are used to combine values coming from different directories.

For example, it is useful to combine entries from LDAP directory with a standard directory provided by Nuxeo5.

18.5.1. Multi-directory sources

```
<component name="org.nuxeo.ecm.directory.multi.config">

  <extension
    target="org.nuxeo.ecm.directory.multi.MultiDirectoryFactory"
    point="directories">

    <directory name="multi">
      <schema>schema</schema>
      <idField>uid</idField>
      <passwordField>password</passwordField>

      <source name="sourceA" creation="true">
        ...
      </source>

      <source name="sourceB">
        ....
      </source>
    </directory>

  </extension>

</component>
```

18.5.2. Sub-directories

```
      <source name="source1" creation="true">
        <subDirectory name="dir1">
          <field for="thefoo">foo</field>
        </subDirectory>
        <subDirectory name="dir2">
          <field for="uid">id</field>
          <field for="thebar">bar</field>
        </subDirectory>
      </source>
```

18.6. The Directory API

This component provides tools to dialog with directories, make queries and get informations. Three classes are a

18.7. Building custom option lists in forms with vocabularies

Directory structure are alse useful to build and store vocabularies which will be used in option lists.

A new vocabulary is declared as a directory:

```
<component name="org.nuxeo.ecm.directories">

  <extension target="org.nuxeo.ecm.directory.sql.SQLDirectoryFactory"
    point="directories">
    <directory name="subject">
      <schema>vocabulary</schema>
      <dataSource>java:/nxsqldirectory</dataSource>
      <cacheTimeout>3600</cacheTimeout>
      <cacheMaxSize>1000</cacheMaxSize>
      <table>subject</table>
      <idField>id</idField>
      <autoincrementIdField>false</autoincrementIdField>
      <dataFile>directories/subject.csv</dataFile>
      <createTablePolicy>on_missing_columns</createTablePolicy>
    </directory>
  </extension>
```

```
</component>
```

The different attributes have the same behaviour as other directories.

Let's have a look at the schema attribute which can take two different values:

- `vocabulary`: this schema is provided to make default vocabulary
- `xvocabulary`: this schema is used to define linked vocabularies: when using `xvocabulary` schema, an other attribute should be defined: `parentDirectory` points the parent directory name to which the current one is relative**.**

Chapter 19. Mimetype detection

19.1. Introduction

The `org.nuxeo.ecm.platform.mimetype.*` packages give all the tools to find the mimetype of a document. The package provides two guessing approach: using file extensions and using the guessing library Jmimemagic (third party tool providing been enhanced detection methods, based on the binary signature of files).

19.2. MimetypeRegistry

All the recognized mimetypes are stored in the `MimetypeRegistry`. Each mimetype definition is a contribution to the `mimetype` extension point of the `org.nuxeo.ecm.platform.mimetype.service.MimetypeRegistryService` component.

```
<component
  name="org.nuxeo.ecm.platform.mimetype.service.MimetypeRegistryService">
  <extension
    target="org.nuxeo.ecm.platform.mimetype.service.MimetypeRegistryService"
    point="mimetype">
    <mimetype normalized="application/vnd.oasis.opendocument.text"
      binary="true" iconPath="odt.png" oleSupported="true">
      <mimetypes>
        <mimetype>application/vnd.oasis.opendocument.text</mimetype>
      </mimetypes>
      <extensions>
        <extension>odt</extension>
      </extensions>
    </mimetype>
  </extension>
</component>
```

A `mimetype` node, bound to a `MimetypeDescriptor` defines a normalized mimetype with the following informations:

- `normalized`: the mimetype entry that is described and that will be returned
- `binary`: a boolean that indicates if the file is a binary one
- `iconPath`: the filename of the image representing the icon
- `oleSupported`: this file mimetype is supported by the `oleExtract` transform plugin - default is `False`
- `onlineEditable`: this mimetype is supported by `online Edit` - default is `False`
- `mimetypes`: list of mimetypes bound to this normalized mimetype
- `extensions`: list of extensions bound to this normalized mimetype

An other defined extension point is `extension` that allow to register extensions that are ambiguous to force mimetype sniffing.

```
<component
  name="org.nuxeo.ecm.platform.mimetype.service.MimetypeRegistryService">
  <extension
    target="org.nuxeo.ecm.platform.mimetype.service.MimetypeRegistryService"
    point="extension">
    <fileExtension name="xml" mimetype="text/xml" ambiguous="true" />
  </extension>
</component>
```

A `fileExtension` node, bound to `ExtensionDescriptor`, has the following attributes:

- `name`: the file extension

- `mimetype`: the associated mimetype
- `ambiguous`: force the mimetype sniffing

At the `MimetypeRegistry` level, methods are provided to dynamically register (or unregister) mimetypes or extensions by code

```
public void testMimetypeRegistration() {
    MimetypeEntry mimetype = getMimetypeSample();
    mimetypeRegistry.registerMimetype(mimetype);
    assertEquals(
            mimetypeRegistry.getMimetypeEntryByName(mimetype.getNormalized()),
            mimetype);
}

private static MimetypeEntryImpl getMimetypeSample() {

    String normalizedMimetype = "application/msword";

    List<String> mimetypes = new ArrayList<String>();
    mimetypes.add("application/msword");
    // fake
    mimetypes.add("app/whatever-word");

    List<String> extensions = new ArrayList<String>();
    extensions.add("doc");
    extensions.add("xml");

    String iconPath = "icons/doc.png";

    boolean binary = true;
    boolean onlineEditable = true;
    boolean oleSupported = true;

    return new MimetypeEntryImpl(normalizedMimetype, mimetypes, extensions,
            iconPath, binary, onlineEditable, oleSupported);
}
```

The `registerMimetype` method, with a `MimetypeEntry` argument, adds the given entry to the registry. After adding the mimetype, one can see that the `MimetypeEntry` retrieved using the initial Normalized name is correct. The `getMimetypeSample` method, helper for the `TestMimetyepRegistryService` test class, shows a definition of a `MimetypeEntry` by code.

Based on the previous entry definition, two registries are built allowing to retrieve the normalized mimetype through file extension (efficient) or by sniffing (accurate).

19.3. Mimetype sniffing

This section describes the underlying process involved in mimetype detections.

The file extension detection rely on the filename that has to be correct. As it is a simple string to associate to the mimetype, we do not develop further. Just note that the filename has to be correctly named and have an extension to retrieve a mimetype with this method.

The sniffing tries to analyse the file content itself in order to guess the mimetype. The third-party tool `Jmimemagic` is used for this and widely enhanced on defining new supported mimetypes and detectors. The `Jmimemagic` tool uses 2 files, `magic.xml` and `magic_1_0.dtd`. We redefine the `XML` one to add new detections (no extension-point there defined from `Jmimemagic` too ;-) l).

Basically, the default mimetype sniffing is based on searching a sequence of characters (or binary values) at a specified `offset`.

```
<match>
  <mimetype>application/pdf</mimetype>
  <extension>pdf</extension>
  <description>PDF document</description>
  <test offset="0" type="string" comparator="=">%PDF-</test>
</match>
```

A `match` is the definition of a magic entry. It contains a `mimetype`, an `extension` and a textual `description` of the defined mimetype. A `test` node, containing the operation to perform is also defined. Here it declares that for an `application/pdf` mimetype, the file has to contain the string `%PDF-` at `offset 0`.

if this method is usually suitable for a lot of files (i.e. one can find some invariants in the format), when used with more complex ones, a simple offset (or a combination) is not enough and we have to refine the detection algorithm. That is what detectors are made for and we have defined some for the 2 major office file formats, MsOffice and OpenOffice.org.

For OpenOffice.org, the zip detection is enhanced

```
<match>
  <mimetype>application/zip</mimetype>
  <extension>zip</extension>
  <description>Zip archive data</description>
  <test offset="0" type="string" comparator="=">PK\003\004</test>
  <match-list>
  <!--  opendocument & OOo 1.x -->
    <match>
      <mimetype>OOo</mimetype>
      <extension>OOo</extension>
      <description>OOo 1.x and OpenDocument file</description>
      <test type="detector" offset="0" length="" bitmask="" comparator="=">
           org.nuxeo.ecm.platform.mimetype.detectors.OOoMimetypeSniffer
      </test>
    </match>
</match>
```

First a simple `offset` detection is performed to qualify a `zip` file, then a sub-match is defined. The detector type indicates that the `org.nuxeo.ecm.platform.mimetype.detectors.OOoMimetypeSniffer` has to be called and this is its responsibility to give all the valid information (`mimetype`, `extension`, `description`) if the file is of correct type.

For MS-Office files, the "magic numbers" (the value to be found at a certain offset) are not that clear, as the magic number defined by `Jmimemagic` (based on the `file` Linux command resource file) is the same for all MS-Office application. Then we invoke a detector for each component that uses the `POI` library to detect what file we deal with.

```
<match>
  <mimetype>application/msword</mimetype>
  <extension>doc</extension>
  <description>Microsoft Office Document</description>
  <test offset="0" type="string" comparator="=">\320\317\021\340\241\261</test>
  <match-list>
    <!--  XLS file by detector -->
    <match>
      <mimetype>application/vnd.ms-excel</mimetype>
      <extension>xls</extension>
      <description>Excel File</description>
      <test type="detector" offset="0" length="" bitmask="" comparator="=">
           org.nuxeo.ecm.platform.mimetype.detectors.XlsMimetypeSniffer
      </test>
    </match>
   <!-- PPT file by detector -->
    <match>
      <mimetype>application/vnd.ms-powerpoint</mimetype>
      <extension>ppt</extension>
      <description>Powerpoint File</description>
      <test type="detector" offset="0" length="" bitmask="" comparator="=">
           org.nuxeo.ecm.platform.mimetype.detectors.PptMimetypeSniffer
      </test>
    </match>
  </match-list>
</match>
```

Once a `Microsoft Office Document` has been detected at `offset 0` in the first match, two sub-matches detectors are defined for `application/vnd.ms-excel` (`org.nuxeo.ecm.platform.mimetype.detectors.XlsMimetypeSniffer`) and `application/vnd.ms-powerpoint` (`org.nuxeo.ecm.platform.mimetype.detectors.PptMimetypeSniffer`). If none returns a correct mimetype, the only possibility remains then `application/msword`. (This may be lightly refactored for simplicity in a near future).

A detector is a class that implements `net.sf.jmimemagic.MagicDetector`. The public `process` method has to detect if the file fulfills the condition. If successful, it returns the mimetypes supported by this file. The public methods `getHandledExtensions` and `getHandledTypes` define the `String` arrays that are used by Jmimemagic to build the final match.

19.4. Invoking the mimetype detection

All the detection is based on `MimetypeRegistry` like object. When invoked, the registry is populated with the information that has been exposed previously. The registry implements the interface `org.nuxeo.ecm.platform.mimetype.interfaces.MimetypeRegistry`.

Once available, directly or from a bean, any `File` or `Blob` can be analyzed and the information retrieved like the mimetype name or the supported extensions list.

```
import org.nuxeo.ecm.platform.mimetype.ejb.MimetypeRegistryBean;
...
    private MimetypeRegistryBean mimetypeRegistry;
...
    public void testSniffWordFromFile() throws Exception {

        File file = FileUtils.getResourceFileFromContext("test-data/hello.doc");

        String mimetype = mimetypeRegistry.getMimetypeFromFile(file);
        assertEquals("application/msword", mimetype);

        List<String> extensions = mimetypeRegistry.getExtensionsFromMimetypeName(mimetype);
        assertTrue(extensions.contains("doc"));
    }
```

In the above example, the `mimetypeRegistry` object used is a `MimetypeRegistryBean`. The purpose is to sniff the mimetype of a given file. The MS-Word file is first read and the `getMimetypeFromFile` method is called. Once the mimetype is retrieved, the `getExtensionsFromMimetypeName` can be called and it gives the associated extensions from the registry.

Note that the API of `org.nuxeo.ecm.platform.mimetype.interfaces.MimetypeRegistry` contains various ways to ask for a mimetype, dealing with `File` or `Blob` objects, with or without default responses. It is worth a look to avoid unneeded work.

Chapter 20. Content Transformation

20.1. Introduction

Transforms are operations that perform any action modifying an input document. This can cover file conversion as well as mail merging or information extraction.

20.2. Plugins module

The various plugins are stored in the `nuxeo-plateform-transforms-plugins.*` module. They are declared as a contribution of the `org.nuxeo.ecm.platform.transform.service.TransformService` component at extension point `plugins`.

The plugins are the engines that perform the needed transformations.

20.2.1. Creating a plugin

A transform plugin follows a framework defined in the `nuxeo-plateform-transforms-core` module. It is basically a class that extends a `org.nuxeo.ecm.platform.transform.interfaces.Plugin`. The main common method that has to be exposed is `transform`

```
    List<TransformDocument> transform(
            Map<String, Serializable> options, TransformDocument... sources);
```

The `transform` method returns a list of `TransformDocument` and accepts an `options map` and `TransformDocument sources`. `TransformDocument` object is defined in `org.nuxeo.ecm.platform.transform.interfaces.TransformDocument` and holds the binary as well as other information such as mimetype.

The `options` map holds all the necessary options to be passed to the plugin. Please note that the keys are the plugin names. See the `officeMerger` plugin below for an implementation example.

20.2.2. Declaring a plugin module

We first declare a contribution to the extension point `plugins` of `org.nuxeo.ecm.platform.transform.service.TransformService`:

```
<extension target="org.nuxeo.ecm.platform.transform.service.TransformService" point="plugins">
  <plugin name="any2odt"
      class="org.nuxeo.ecm.platform.transform.plugin.joooconverter.impl.JOOoConvertPluginImpl"
      destinationMimeType="application/vnd.oasis.opendocument.text">
    <sourceMimeType>text/xml</sourceMimeType>
    <sourceMimeType>text/plain</sourceMimeType>
    <sourceMimeType>text/rtf</sourceMimeType>

    <!-- Microsoft office documents -->
    <sourceMimeType>application/msword</sourceMimeType>

    <!-- OpenOffice.org 1.x documents -->
    <sourceMimeType>application/vnd.sun.xml.writer</sourceMimeType>
    <sourceMimeType>application/vnd.sun.xml.writer.template</sourceMimeType>

    <!-- OpenOffice.org 2.x documents -->
    <sourceMimeType>application/vnd.oasis.opendocument.text</sourceMimeType>
    <sourceMimeType>application/vnd.oasis.opendocument.text-template</sourceMimeType>

    <option name="ooo_host_name">localhost</option>
    <option name="ooo_host_port">8100</option>
  </plugin>
</extension>
```

The name attribute will be used to declare the transform chain. The class attribute is the class that will do the effective job of the transformation while the destinationMimeType is the mime-type of the result of the transform.

After attributes, <sourceMimeType> nodes define the allowed input mime-types the transform is supporting. In the presented case, we can see that the any2odt plugin will be able to handle text, Microsoft office word , OOo 1.x and OpenDocument format (OOo2.x) files to output an application/vnd.oasis.opendocument.text OpenDocument file. Options can also be added as we see for <option> tags with ooo_host_name and ooo_host_port attributes.

Plugins can be combined to build transform chains. This chains are declared in a transformer which is a contribution to the extension-point transformers of org.nuxeo.ecm.platform.transform.service.TransformService component.

```
<extension target="org.nuxeo.ecm.platform.transform.service.TransformService"
           point="transformers">
  <transformer name="any2text"
               class="org.nuxeo.ecm.platform.transform.transformer.TransformerImpl">
    <plugins>
      <plugin name="any2pdf"/>
      <plugin name="pdf2text"/>
    </plugins>
  </transformer>
</extension>
```

A transformer is defined by its name. This is this name that will be used to initialize a transform service when using it.

Then, plugins involved in the chain are listed. Wen can see that our any2txt transformer is composed with two chained plugins: any2pdf then pdf2txt. Obviously, a single plugin for a transform is legal as we can see with the use of our previous any2odt plugins.

```
<extension target="org.nuxeo.ecm.platform.transform.service.TransformService"
           point="transformers">
  <!-- This transformer uses a the OOo plugin to transform documents to ODT-->
  <transformer name="any2odt"
               class="org.nuxeo.ecm.platform.transform.transformer.TransformerImpl">
    <plugins>
      <plugin name="any2odt"/>
    </plugins>
  </transformer>
</extension>
```

20.2.3. Using a transform plugin

Once a transform plugin has been declared and the transformer is known, we can use it to perform various transformation actions.

```
TransformService service = NXTransform.getTransformService();
Transformer transformer = service.getTransformerByName("any2pdf");

List<TransformDocument> results = transformer.transform(null,
        new TransformDocumentImpl(sourceStream, "application/vnd.oasis.opendocument.text"));

SerializableInputStream resultStream = results.get(0).getStream();
```

We first get the TransformService that exposes all the available transforms. Then a specific transformer is built with the getTransformerByName method of the service. The name is the one that has been declared in the contribution to the transformers extension-point of the org.nuxeo.ecm.platform.transform.service.TransformService component.

Then the transformer exposes a transform method that returns a list of TransformDocument. The arguments are the:

- options: the plugin options (the keys are the plugin names) - we pass null here as we do not have any

options. See `officeMerger` plugin for a more detailed example

- list of sources as `TransformDocument` instances

There are three levels of options that overload. First the `plugin` options are the default. Then any option in the `transformer` that define an option for this plugin overload them. Finally, any code-defined options are merged, overloading any previous option that may have been already defined. Please note again that options are defined on a plugin name basis.

A `TransformDocument` instance can be constructed from:

- a source stream
- a source stream and its mime-type
- a blob

In our example, we use the second way and give a `sourceStream` and the mime-type of an ODF document.

Once the input list processed, the results list contains all the transformed files as `TransformDocument` instances from which you can retrieve the `SerializableInputStream` stream with `getStream` method, the mime-type using `getMimetype` and the blob with `getBlob` method.

An alternate way of using the transform is to call it directly from the service

```
List<TransformDocument> results = service.transform(converter,
        null, new TransformDocumentImpl(stream, sourceMimetype));
```

The arguments are the same and the `converter` name is given as first argument. An other constructor using blobs as input instead of `TransformDocument` is also available. Before calling a transform we can also check that the source mime-type is supported by calling the `isMimetypeSupportedByPlugin` method. Be careful though that this plugin name may be different than the transformer name.

Transforms can be called directly but are also part of the docModifier framework that reacting on events, can call transforms to alter or generate new informations (see below oleExtract plugin or docModifier documentation)

20.3. Available transforms

All the transform plugin packages start with `org.nuxeo.ecm.platform.transform.plugin`. Some of them are optional and not included in the default platform.

20.3.1. Document conversion

The document conversion plugin is a generic transform delivered in Nuxeo that allow transforming a file from a format to an other. It uses the third-party JODconverter tool. The transform is implemented in the `org.nuxeo.ecm.platform.transform.plugin.joooconverter.impl.JOOoConvertPluginImpl` class and defined as a `plugin` contribution to `org.nuxeo.ecm.platform.transform.service.TransformService` .

Then it can be classically called as explained in the previous section:

```
TransformService service = NXTransform.getTransformService();
Transformer transformer = service.getTransformerByName("any2pdf");

List<TransformDocument> results = transformer.transform(null,
        new TransformDocumentImpl(sourceStream, "application/vnd.oasis.opendocument.text"));

SerializableInputStream resultStream = results.get(0).getStream();
```

The `transform` call relays to the `JOOoConvertPluginImpl` that first connects to OOo on port and host defined in the contribution (usually `localhost:8100`) or in the options (not defined in our example). It then acquire an `OpenOfficeDocumentConverter` from JODconverter tool and then can call the `convert` method with the requested target mimetype and source file. Note that in future version, the `StreamOpenOfficeDocumentConverter` class will be used to avoid dealing with `File` objects. This limitation will be solved when a new version of OpenOffice.org (2.3) will be out and solves a regression on loading streams.

Before any call to the underlying OpenOffice.org converter, once in the transformation engine `JOOoConvertPluginImpl`, the source document mimetype is tested. If it is the same than the requested `destinationMimetype` defined by the plugin, the source file is returned immediately as result unless the mimetype occurs in the `<sourceMimeType>` of the plugin to allow self-transformations. By default, `any2pdf` plugin will return immediately if an `application/pdf` file is submitted while any OpenDocument transformation (`any2odt`, `any2ods`, `any2odp`) will process the files as each one contains its own mimetype like `<sourceMimeType>application/vnd.oasis.opendocument.text</sourceMimeType>`. This allow to clean and validate files forged by hand and possibly apply automatic treatments at OpenOffice.org side.

20.3.1.1. Remote engine

Since OOo 2.3.0, it is possible to use `InputStreams` as source documents so that no more `File` objects are needed. This is useful if we want to isolate the OOo server on a separate machine (and if needed use a farm with load balancing if heavy work is needed) As this feature is only available with OOo2.3+ (due to a bug in previous versions), there is a `nuxeo.property` to define the OOo version used.

```
org.nuxeo.ecm.platform.transform.ooo.version=2.3.0
```

At term, the JODConverter tool should be able to return the OOo version it is using, so that this property will not be needed anymore. For the moment, let use this property. But as OOo stream loading has been reported as less efficient than `File`/`URL` by JODConverter users, the streaming method is only used if OOo is really a remote instance. This is analyzed if the `ooo_host_name` transformer option is '`localhost`' or starts with '`127.0.0`'. And of course, only if OOo declared version is greater than 2.3.0.

20.3.1.2. OpenOffice.org loading options

As the underlying engine is based on OpenOffice.org, one can extend the document converter by supporting OOo loading options. The transform plugins options mechanism is fully supported, so they can be defined at plugin, transformer and code level. The available fields to be passed to OOo is listed at MediaDescriptor IDL reference and have to be bound to the plugin name.

Here is an example passing two options for loading conditions of a special document that is a password protected and has an autostart bound event that delete the content of the file to put the word DELETED in the document. This simulate any other heavy process such as document merging or automatic mail-merge from a database.

```
Map<String, Map<String, Serializable>> options = new HashMap<String, Map<String, Serializable>>();
Map<String, Serializable> pdfOptions = new HashMap<String, Serializable>();

pdfOptions.put("Password", "TheDocumentPassWord");
short ALWAYS_EXECUTE_NO_WARN = 4;
pdfOptions.put("MacroExecutionMode", ALWAYS_EXECUTE_NO_WARN);
pdfOptions.put("ReadOnly", false);
pdfOptions.put("Hidden", false);
options.put("any2pdf", pdfOptions);

// Note: due to password, file mimetype can not be sniffed
// so mimetype has to supplied to TransformDocumentImpl
results = service.transform("any2pdf", options,
        new TransformDocumentImpl(getBlobFromPath(path), "application/vnd.oasis.opendocument.text"));
assertTrue(results.size() > 0);

// The macro on the open event replaces the text by "DELETED"
assertEquals("pdf content", "DELETED", DocumentTestUtils.readPdfText(pdfFile));
```

First, we build the plugin options map. The `Password` field is in charge of sending the password string to the OOo loader so that the file can be opened. The `MacroExecutionMode` field defines how macros and script are

handled at startup. By default, the consequence of the `-headless` OOo mode, is that the `NEVER_EXECUTE = 0` value is used. One can change this default behaviour by using any value listed in the OOo MacroExecMode constant group. One important thing to be noted is that type is important: `Password` require a `String` argument while `MacroExecutionMode` requires a `short` one. The types have to be correct otherwise the field will not be handled by OOo.

By default, JODConverter sets the `ReadOnly` option as true. As we want to modify the document, we will have to set the `ReadOnly` flag to `false`. The `Hidden` flag has also to be set to false. This trick is due to a problem in OpenOffice.org PDF engine that seems not to be able to handle document modification while `Hidden`. As an headless server-deployed OOo instance, this should not be a major problem.

Once Options have been defined, they are globally merged to the options map under the plugin name key (here, `any2pdf`)

Then the transform can be called as usual. Please note that mimetype of password document have to be passed explicitly to the `TransformDocumentImpl` constructor as it can not be sniffed by the `Mimetype` service for the moment.

Finally, as expected a PDF document is returned with its content changed to the `DELETED` string.

20.3.2. Pdfbox

20.3.3. OLE objects extraction

OLE objects are objects included in office files that can be edited as standalone ones. For example, a spreadsheet table may be included in a report so that the presented datas are always up to date.

This plugin is located in `nuxeo-plateform-transform-plugin-oleextract` module and is not include by default in the plateform.

20.3.3.1. Implementation

The purpose of this plugin is to extract all these Ole objects and provide them as standalone files so that they can be checked individually. It is also extended to extract images. It has a classical Transform plugin structure, the plugin name is `oleExtractPlugin` bound to `org.nuxeo.ecm.platform.transform.plugin.oleextract.impl.OfficeOleExtractorPluginImpl`. The transform name is `oleExtract`.

As an example it can be called like

```
List <TransformDocument > results =
        service.transform(TRANSFORMER_NAME, null, new TransformDocumentImpl(stream, mimetype));
TransformDocument result = results.get(0);
List <Map<String, Serializable>> ole =
        (List < Map < String, Serializable > >) result.getPropertyValue("ole:olecontents");
```

First, the transform service is classically called, the `TRANSFORMER_NAME` being set to `oleExtract`. After processing, the only `TransformDocument` returned result contains a property `ole:olecontents` that gives the list of embedded objects that have been extracted.

The `olecontents` schema is defined in `olecontent.xsd`. Each element of the list contains the following fields

```
<xs:complexType name="olecontent">
  <xs:sequence>
    <xs:element name="displayname" type="xs:string" />
    <xs:element name="filename" type="xs:string" />
    <xs:element name="mime-type" type="xs:string" />
    <xs:element name="data" type="nxs:content"/>
    <xs:element name="thumbnail-mime-type" type="xs:string"/>
    <xs:element name="thumbnail-data" type="nxs:content"/>
  </xs:sequence>
</xs:complexType>
```

Each `olecontent` element contains file datas (`data` & `mime-type`) and thumbnail ones (`thumbnail-mime-type` & `thumbnail-data`). The `displayname` is the name retrieved in the office file if it was named otherwise the internal one that has been given in the office file. The `filename` fields is built from the `displayname` and the extension deduced from the `mime-type`.

Based on this new schema, the File doctype is extended in the org.nuxeo.ecm.platform.transform.oleextract.coretypes contributing type and schema needed by oleExtract

```
<extension target="org.nuxeo.ecm.core.schema.TypeService" point="doctype">
  <doctype name="FileWithOle" extends="File">
    <schema name="olecontents" />
  </doctype>
</extension>
```

A new `FileWithOle doctype`, based on `File`, is defined. It can be subtype of `Workspace` and `Folder` and is defined as a contribution of `org.nuxeo.ecm.platform.types.TypeService` which allows to create new documents based on it.

The oleExtract transform plugin has been bound to the `oleExtractModifier` docModifier. The contribution to the extension point is defined in

```
<extension target="org.nuxeo.ecm.platform.modifier.service.DocModifierService"
           point="docTypeToTransformer">
  <documentation>
    docModifier for oleExtract transform plugin.
  </documentation>
  <docModifier name="oleExtractModifier"
               documentType="FileWithOle"
               transformationPluginName="oleExtract"
               sourceFieldName="file:content"
               destinationFieldName="file:content">
    <coreEvent>documentCreated</coreEvent>
    <coreEvent>documentModified</coreEvent>
    <customField name="olecontents:olecontents" transformParamName="ole:olecontents"/>
    <customOutputField outputParamName="ole:olecontents" name="olecontents:olecontents"/>
  </docModifier>
</extension>
```

This contribution reacts on document creation and modification. It receives the initial office file `file:content` and gives back the `olecontents:olecontents` back mapped top the `ole:olecontents` we saw above. The initial file is returned unchanged.

With this docModifier reacting on some events, the oleExtract results can now be integrated in the application. The org.nuxeo.ecm.platform.transform.oleextract.action component defines an ActionService contribution

```
<action id="TAB_OLEOBJECT" link="/incl/tabs/document_oleobjects.xhtml"
        enabled="true" label="action.view.ole" order="49">
  <category>VIEW_ACTION_LIST</category>
  <filter id="view_ole">
    <rule grant="true">
      <type>FileWithOle</type>
    </rule>
  </filter>
</action>
```

The `TAB_OLEOBJECT` action defines a new tab listing the `olecontents:olecontents` elements and providing links for retrieving each file provided the document is of the current type `FileWithOle`.

20.3.3.2. Extraction details

OleExtract is based on the parsing of OpenDocument File format. If the submitted file is an ODF one, then it is unzipped and processed without any connection to OpenOffice.org. If the file is not an ODF one, then a converter plugin is used according to the source mime type. In this case, OpenOffice.org is required as a JODConverter dependency.

Once we have an ODF file it is unzipped and its content.xml is parsed to find `draw:object-ole`, `draw:object`

and draw:image related tags. For each one, the name of the resource is retrieved if it exists. Then for each found resource, the alternate view is retrieved so that a preview can be proposed when listing this content (still under development)

ODF resources in an ODF file (think at a spreadsheet diagram embedded in a text document) are stored in directories and flat XML form while other resources are stored in a binary format. So for ODF resources the global manifest file is parsed to isolate the files with their correct manifest:media-type so that the new ODF archive for the Ole object can be built. Once the new manifest file is created, the embedded ODF directory is zipped and this binary form is returned.

20.3.4. Office files merger

This transform plugin is contained in nuxeo-plateform-transform-plugin-officemerger module and is not include by default in the platform. Its purpose is to build a new file as the result of the merging of the list given as parameters. It uses OpenOffice.org merging capabilities, mainly through the insertDocumentFromURL UNO method for text documents.

The public method merge is available from OfficeMergerImpl and returns a SerializableInputStream containing the resulting document.

```
OfficeMergerImpl merger = new OfficeMergerImpl();
SerializableInputStream result = merger.merge(sourceFiles, engineType, converter,
         outlineRank, withPageBreaks);
SerializableInputStream result = merger.mergeStreams (sourceStreams, engineType,
         converter, outlineRank, withPageBreaks);
```

Some options have been added to enhance the building of the main document. Here is the list of the arguments of the merge method

- sourceFiles/sourceStreams: Ordered array of File objects or streams to be merged.
- engineType: Depending on the nature of source documents, the OpenOffice.org API to be used is obviously not the same. If source files are text file, the user will probably want to have a text file as a result while if he deals with slides, the results is expected as a presentation. This String argument tells which engine to be used (text, presentation, spreadsheet- only text is already implemented) - Default text
- converter: Once the document is built, the Document Converter plugin can be called automatically to create the final document. This is the converter name that is expected (eg. any2pdf) and an exception is raised if it does not exist or the mime type deduced from the engineType is not supported. If an empty string is provided, no transform occur at the end of the merging - Default empty
- outlineRank: This is the rank (compared to the file list) where a Table Of Content may appear. For example, if the value is 3, then the first two files of the file list are inserted, the T.O.C is built and inserted and then the remaining files are processed. The Table of Content is refreshed at the end of the whole insertion. A value of 0 means no T.O.C. - Default 0
- withPageBreaks: A boolean that adds or removes page breaks between file insertions - Default true

The plugin engine, the merge method is available directly but the principal use will occur through a Transform call. The Transform name is officeMerger and the package name is org.nuxeo.ecm.platform.transform.plugin.officemerger

```
<plugin name="OfficeMergerPlugin"
        class="org.nuxeo.ecm.platform.transform.plugin.officemerger.impl.OfficeMergerImpl"
        destinationMimeType="application/vnd.oasis.opendocument.text">
  <sourceMimeType>application/msword</sourceMimeType>
  <sourceMimeType>application/vnd.oasis.opendocument.text</sourceMimeType>
  <sourceMimeType>application/vnd.sun.xml.writer</sourceMimeType>
  <option name="ooo_host_name">localhost</option>
  <option name="ooo_host_port">8100</option>
</plugin>
```

Note that OpenOffice.org has to be listening from incoming UNO connections on the specified interface

ooo_host_name and port ooo_host_port. If it is not the case, an OpenOfficeException exception will be raised.

So, once defined, the officeMerger transform can be called passing the options like

```
Map<String, Serializable> mergingOptions = new HashMap<String, Serializable>();
mergingOptions.put("engineType", "text");
mergingOptions.put("converter", "any2pdf");
mergingOptions.put("outlineRank", 0);
mergingOptions.put("withPageBreaks", false);

options.put("officeMerger", mergingOptions);

List<TransformDocument> results = transformer.transform(options, sourceFiles);
```

Note that mergingOptions can be incomplete or even null. The options will then take their default values.

The results list contains the final merged document, and converted if requested, at first index.

20.3.5. XSL Transformation

The XSLT plugin provides XSL transformations in Nuxeo. It allows you to transform XML documents using a XSL stylesheet as defined in the XSLT Specification. The plugin is implemented in the org.nuxeo.ecm.platform.transform.plugin.xslt.impl.XSLTPluginImpl class and defined as a plugin contribution to org.nuxeo.ecm.platform.transform.service.TransformService.

The XSLT Plugin accepts XML documents as source files, and you must provide the XSL stylesheet as a plugin's option named stylesheet. The XSL stylesheet must be provided as a Blob.

Then, you can easily transform your documents:

```
final Map<String, Serializable> pluginOptions =
        new HashMap<String, Serializable>();
pluginOptions.put("stylesheet", (FileBlob) getXSLStylesheetBlob(...);
final Map<String, Map<String, Serializable>> options =
        new HashMap<String, Map<String, Serializable>>();
options.put("xslt", pluginOptions);

TransformServiceCommon service = TransformServiceDelegate.getLocalTransformService();
final List<TransformDocument> results = service.transform(
        "xslt", options, xmlSourceFiles);
```

The resulting documents' mime-type is set depending of the method attribute of the xsl:output element, specified in the XSL stylesheet. If there is no method attribute defined in the XSL stylesheet, a default value is chosen for the method attribute as defined in the XSLT Specification.

The mime-type is chosen as follows:

- text/html: if the output method is html.
- text/plain: if the output method is text.
- text/xml: if the output method is xml.

If there is no XSL stylesheet provided to the plugin or if an error occurs during the transformation (corrupted xml or xsl for instance), a TransformException is thrown.

Chapter 21. Nuxeo Conversion Service

The `nuxeo-core-convert` provides a service to manage conversion of Blobs from one format to an other.

`nuxeo-core-convert` is only available starting with 5.2-M4.

21.1. Conversion Service vs Transformation Service

In 5.2 Conversion Service replaces the Transformation service that is now deprecated.

21.1.1. Motivations for this API changes

Transformation service had some API design issues that we wanted to correct. Because full text indexing is now handled by the repository we also had to have a core service to manage full text conversion. We decided to define a brand new API with a new service and we changed the service name to avoid any confusion and be able to provide backward compatibility.

21.1.2. What has been improved

- API is now simpler

  there are only converters, not transformers and plugins like before

- ConversionService includes a caching system

  this eliminate the need of having custom cache managed by all high level services that may use converters (like the preview service)

- Data input/output is now handled via `BlobHolder` interface.

  There is only data structure (no more TransformDocuments or plain Blobs). This also makes the caching system more efficient sice link between the blobs and the associated `DocumentModel` is preserved when available.

- Availability check API interface.

  ConversionService now provides an API to know if converter is available, this is usefull when the converter depends on an external program that must be installed on the server (like OpenOffice server)

21.1.3. About compatibility

Transformer API is now deprecated and the old transform-service implementations and plugins have been removed from default distribution. Nevertheless, we provide a tranformer-compat bundle that handles compatibity between the old and the new API.

In order to activate this compatibility you need to deploy :

- `nuxeo-platform-transform-api`
- `nuxeo-platform-transform-compat`

21.1.3.1. Using transformer API

Transformers API is stil available, but the implementation now wraps calls to the ConversionService. This means you can use the new Converters from the old Tranformation service API. All default included

Transformers have been migrated to converters with the same name. Code that was using transformers should still work in 5.2.

21.1.3.2. Contributed transformers and plugins

Contributions to the old Transformation service are now contributed to the ConversionService using specific converters that wraps transformers or Plugins.

There are some limitations thought :

- Transformers must be based on the TransformerImpl class
- you may have to change the dependencies in pom.xml and MANIFEST.MF to point to the compat artifact instead of the core one.

21.2. Using Conversion Service

21.2.1. About `BlobHolder`

BlobHolder defined an abstract interface for objects that can hold one or several blobs. BlobHolders can be created from plain `Blob` (or `List<Blob>`) or from a `DocumentModel`. Compared to a plain `Blob`, a `BlobHolder` provides lazy loading of a blob when possible and keeps track of the link between the Blob and the DocumentModel. In fact, BlobHolder interface is also available as a DocumentModelAdapter, so you can simply get a BlobHolder from a DocumentModel. Using BlobHolder provides several advantages :

- no more need to hardcode the xpath to extract the blob from a DocumentModel

  (this is done once for all in the DocumentModelAdpater configuration)

- conversionservice caching system can no efficiently monitor changes

```
# create a BlobHolder from a File

File file = ....

Blob blob = new FileBlob(file);
BlobHolder fileBlobHolder = new SimpleBlobHolder(blob);

# get BlobHolder from a Document

DocumentModel myDoc = ...

BlobHolder docBlobHolder = myDoc.getAdapter(BlobHolder.class);


# create a BlobHolder "by hand" on top of a DocumentModel

BlobHolder docBlobHolder2 = new DocumentBlobHolder(myDoc, "file:content");
```

21.2.2. built-in converters

Inside `nuxeo-core-convert-plugins`

- pdf2text, xml2text, html2text, word2text, xl2text, ppt2text, oo2text

  text extractors for common office formats

- rfc822totext

text extractors mime encoded mails

- any2text

  meta-converter for text extraction

Inside `nuxeo-platform-convert`

- pdf2html

  PDF to html conversion based on pdftohtml command line tool

- office2html

  convert standard office formats to html (uses openoffice)

- any2html

  compound converted to convert any input to html

- any2pdf

  cuse OpenOffice to generate PDF

21.2.3. Conversion Service API

Conversion Service can be accessed via the standard Nuxeo Service lookup :

```
ConversionService conversionService = Framework.getService(ConversionService.class);
```

To convert a `BlobHolder` to a given destination mime type :

```
BlobHolder result = conversionService.convertToMimeType("text/plain", blobHolder, params);
```

params is a simple `Map<String,Serializable>` to pass parameters to the converter (can be null);

To use a known converter :

```
BlobHolder result = conversionService.convert("converterName", blobHolder, params);
```

To find a converter to a given conversion :

```
String converterName = conversionService.getConverterName(sourceMimeType, destinationMimeType);
```

To test if a converter is available

```
ConverterCheckResult checkResult = conversionService.isConverterAvailable("converterName");
```

This call can throw `ConverterNotRegistred` if the target converter does not exist at all. The `ConverterCheckResult` class provides :

- a isAvailable() method
- a getErrorMessage() method

  returns the error that occured while doing the availability check

- a getInstallationMessage method

Return the installation message that was contributed by the converter contributor

21.2.4. Configuring converter service

The conversion service supports a global configuration via XML file in order to configure caching.

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.core.convert.config">
  <extension target="org.nuxeo.ecm.core.convert.service.ConversionServiceImpl"
      point="configuration">

    <configuration>
      <!-- define directory location for caching : default to java default tmp dir (java.io.tmpdir) -->
      <cachingDirectory>/var/ConversionCache</cachingDirectory>
      <!-- GC interval in minutes (default = 10 minutes ) -->
      <gcInterval>10</gcInterval>
      <!-- maximum size for disk cache in KB (default to 10*1024) -->
      <diskCacheSize>1024</diskCacheSize>
      <!-- Enables or disables caching (default = true)-->
      <enableCache>true</enableCache>
    </configuration>
  </extension>

</component>
```

21.2.5. Contributing converters

To contribute a converter, you have to contribute a class that implement the `org.nuxeo.ecm.core.convert.extension.Converter` interface. This class will be associated to :

- a converter name
- a list of source mime-types
- one destination mime-type
- optional named parameters

```
<extension target="org.nuxeo.ecm.core.convert.service.ConversionServiceImpl"
    point="converter">

  <converter name="html2text" class="org.nuxeo.ecm.core.convert.plugins.text.extractors.Html2TextConverter">
      <sourceMimeType>text/html</sourceMimeType>
      <sourceMimeType>text/xhtml</sourceMimeType>
      <destinationMimeType>text/plain</destinationMimeType>
      <parameters>
        <parameter name="myParam">myValue</parameter>
      </parameters>
    </converter>

</extension>
```

You can also contribute a converter that is a chain of existing converters (what was called a transformer in 5.1 transform service API). To to this, the contributed transformer does not have to define an implementation class, just a chain of either converters or mime-types. If mime-types are used, the conversion service will automatically guess the converter chain from the mime-types steps.

```
<extension target="org.nuxeo.ecm.core.convert.service.ConversionServiceImpl"
    point="converter">

  <!-- explicit chain of 2 converters : converter1 + converter2 -->
  <converter name="chainedConverter" >
      <sourceMimeType>some/mimetype</sourceMimeType>
      <destinationMimeType>some/other-mimetype</destinationMimeType>
      <conversionSteps>
        <subconverter>converter1</subconverter>
        <subconverter>converter2</subconverter>
      </conversionSteps>
    </converter>

  <!-- define chain via mime types : foo/bar1 => foo/bar2 => foo/bar3 -->
```

```
  <converter name="chainedMimeType" >
      <sourceMimeType>foo/bar1</sourceMimeType>
      <destinationMimeType>foo/bar3</destinationMimeType>
      <conversionSteps>
        <step>foo/bar2</step>
      </conversionSteps>
    </converter>

</extension>
```

When using chained converters, the additionnal optionnal parameters are passed to each underlying converter.

Converter based on external tools (such as command line or OpenOffice server based) can implement the `ExternalConverter` interface. This interfaces adds a isConverterAvailable() method that will be called in order to check converter availability.

21.2.6. Converters based on external command line tools

A lot of conversion tools comes as command line executable. So in some case it's interesting to wraps these command lines into a converter.

For that purpose, we provide a base class for converters that are based on a command line wrapped by the nuxeo commandLine service.

The base class `org.nuxeo.ecm.platform.convert.plugins.CommandLineBasedConverter` handles all the dirty work, and you only have to override the methods to define the parameters of the command line and the parsing of the output.

```
<extension target="org.nuxeo.ecm.core.convert.service.ConversionServiceImpl"
    point="converter">
 <!-- converter based on the pdftohml command line -->
 <converter name="pdf2html" class="org.nuxeo.ecm.platform.convert.plugins.PDF2HtmlConverter">
      <sourceMimeType>application/pdf</sourceMimeType>
      <destinationMimeType>text/html</destinationMimeType>
      <parameters>
        <parameter name="CommandLineName">pdftohtml</parameter>
      </parameters>
    </converter>

</extension>
```

Chapter 22. Relations

22.1. Introduction

Relations in Nuxeo EP 5 follow concepts as described by the W3C Resource Description Framework (RDF).

The purpose is to provide content management relations (relations between documents of the site for instance) as well as being able to share this information with third party applications, by following the RDF standards.

22.2. Concepts

There are a few jargon terms to understand when dealing with relations.

Let's consider a relation like "document A is a version of document B". This relation is described as a *triplet* or *statement*: it has a subject, "document A", a predicate, "is version of", and an object, "document B".

The statement elements are more generally referred to as *nodes*. More specific kinds of nodes are *literals* and *resources*. A subject and a predicate will always be resources, while the object may be also a literal. In a relation like "document A has title 'documentation'", the object will be the literal string 'documentation'.

Literals are simple nodes, holding information like a string or a date. Resources refer to uniquely identifiable objects, and often use a URI as identifier that looks like a URL. If this URI refers to an identified namespace, we can make a difference between resources using it.

For instance, we can use the dcterms namespace to identify predicates: "http://purl.org/dc/terms/References", "http://purl.org/dc/terms/IsBasedOn","label.relation.predicate.IsBasedOn",...

Documents in the Nuxeo default application use the following namespace: "http://www.nuxeo.org/document/uid/". A document URI would look like "http://www.nuxeo.org/document/uid/618e53c8-409e-40e8-9b73-5493f7e6de88" because we use the JackRabbit identifier to identify fully the resource. Imagine that we use custom unique identifiers for documents, we could use them too, but we should use a different namespace so that we do not mistake the JCR identifier and the custom identifier.

When defining a relation like "document A is a version of document B", we will then build a statement which subject is a resource representing document A, which predicate is a resource representing the "is a version of" information, and which object is a resource representing document B.

If we would like to state that this relation was created as a certain date, we will add a date property to the statement. This can be seen as a relation where the subject would be the statement itself, the predicate a resource representing the "was created at" information, and which object would be a literal representing the given date.

22.3. Configuration

If you would only like to change the storage used for the default graph of Nuxeo EP 5, please refer to Section 40.4.2, “Relation service configuration”.

22.3.1. Graph instances

Relations are stored in a *graph*, that can also be called a *model*.

The graph definition is made though an extension point. It holds configuration about where and how to store relations. Here is an example contribution.

Example 22.1. Jena graph configuration for the Relation Service using PostgreSQL as storage

```
<?xml version="1.0" encoding="UTF-8"?>
<component name="MyJenaGraph">
  <require>org.nuxeo.ecm.platform.relations.jena</require>
  <extension target="org.nuxeo.ecm.platform.relations.services.RelationService"
    point="graphs">
    <graph name="default" type="jena">
      <option name="backend">sql</option>
      <option name="databaseType">PostgreSQL</option>
      <option name="datasource">java:/nxrelations-default-jena</option>
      <option name="databaseDoCompressUri">false</option>
      <option name="databaseTransactionEnabled">false</option>
      <namespaces>
        <namespace name="rdf">
          http://www.w3.org/1999/02/22-rdf-syntax-ns#
        </namespace>
        <namespace name="dcterms">http://purl.org/dc/terms/</namespace>
        <namespace name="nuxeo">http://www.nuxeo.org/document/uid/</namespace>
      </namespaces>
    </graph>
  </extension>
</component>
```

This graph uses a Jena graph. Jena is a RDF framework, a plugin has been developed to integrate it to the nuxeo platform. The graph definition requires the plugin to be registered to the application.

The graph is named "default" and declares its connection configuration.

22.3.2. Resource adapters

The graph configuration includes namespaces used for some of the graph resources so that resources with a known namespace can be transformed into any kind of object.

For instance, the namespace "http://www.nuxeo.org/document/uid/" is used to identify documents using their JCR unique identifier. We can register an adapter so that the resource can be transformed into the actual document model it represents.

For example, the DocumentModelResourceAdapter class allows to get a DocumentModel object from a resource, build with a namespace and a local name.

22.4. Manage relations

Managing relations turns around the following actions:

- adding a new relation to a graph
- removing an existing relation from a graph
- determining if a relation already exists
- emptying a graph to remove all existing relations

22.5. Display relations

The StatementInfo interface provides methods to maniulates data about statements. That way, you can retrieve informations to display relations and their properties: These tools are:

- wrappers/getters, to recover information about a statement (subject, predicate, object, ...) or a node (type of node : literal, resource, QNameResource ...)

- methods to get all statements or statements which are matching a given pattern: this can be useful to determine if a relation is incoming (document is the subject of the statement) or outgoing (the document is the object of the resource)
- methods to get extra properties: for each relation, informations are added like creation date or the creator of the relation.

Another interesting interface is NodeInfo: it is useful to qualify the type of resource you are handling with. Some methods can determine if the node is a literal, a document, a blank node, etc ...

22.6. Architecture overview

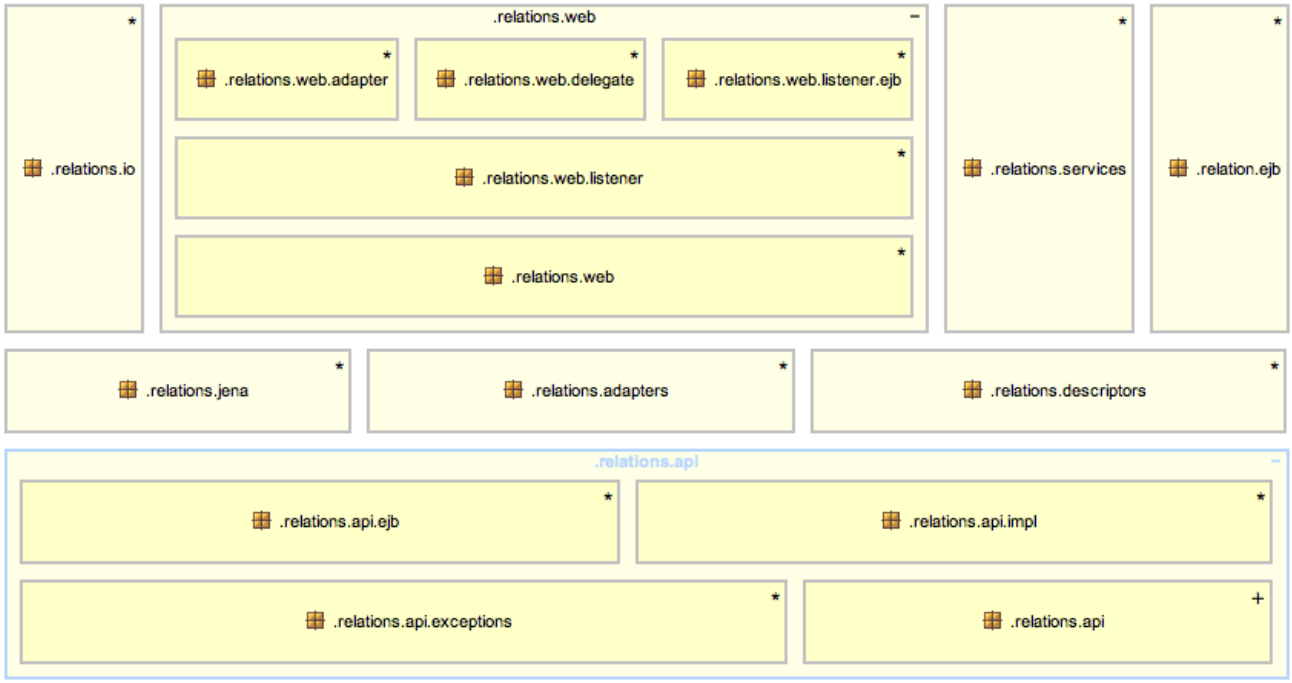

Chapter 23. Placeful Configuration

23.1. Introduction

The placeful configuration service allows configuration to be placed on a node in a repository. It is possible to update, remove this configuration and merged it with all the configuration located from this node to the root of the repository.

Placeful Configuration (PC in the rest of this chapter) is useful when you have a number of nodes much bigger than the number of configuration. It allows to change and merge these configuration without having to travel the repository tree.

23.2. Using Placeful Configuration

The `PlacefulConfigurationManager` is used for all interaction with PC. This service is available via the extension point `org.nuxeo.ecm.platform.placeful.configuration.service.PlacefulConfigurationService`. Before using it, you need to associate a configuration with a storage in the component definition ???.

You can then use the placeful configuration service. The following code snippet show the basic usage:

```
Path p1 = new Path("/mon/path");
RepositoryLocation repo = new RepositoryLocation("monrepo");
NuxeoPrincipal principal = new NuxeoPrincipalImpl("myself");

PlacefulConfigurationManager pcs = Framework.getService(PlacefulConfigurationManager.class);
LocalTheme lt1 = pcs.createConfigurationEntry(LocalTheme.class, p1, repo, principal);
lt1.setMode("myMode");
pcs.saveConfigurationEntry(lt1);

Map<String, Object> map = new HashMap<String, Object>();
map.put("mode", "myMode");
List<LocalTheme> list =  pcs.getAllConfigurations(LocalTheme.class, map);                ❶

LocalTheme lt = pcs.getConfiguration(LocalTheme.class, repo, p, principal);
LocalTheme ltMerge = pcs.getMergedConfiguration(LocalTheme.class, repo, p, principal);
pcs.removeConfiguration(LocalTheme.class, repo, p1, principal);
```

❶ Query the storage for all the PC that have those field/value.

Note that:

- you never create a PlacefulConfiguration yourself but ask the manager for one.
- You need to save the PC after modification.
- You can not move a PC, you need to remove it and create a new one.

For more information on available methods and class, have a look at the Javadoc.

23.3. Contributing a placeful configuration

A PC is composed of two distinct types of information:

- The information specific to this configuration, for example, a local theme configuration has a mode, a page or a docId.
- The information relative to its "placefulness": the path, principal and repository.

To contribute a configuration to the service you only have to gives the information specific to the configuration.

The "placeful" part is taken care of by the service. You also need to give a way to merge the PC. We will create a simple config as an exemple. A "Simple" PC that has only one field: value.

- Create an interface specific for your configuration.

```
interface SimpleConfig {
    setValue();
    getValue();
}
```

- Create the empty interface that will be manipulated by the user. It needs to extends `PlacefulConfiguration`, `Serializable` and your specific interface.

```
public interface Simple extends PlacefulConfiguration, SimpleConfig, Serializable {}
```

- Create the implementation of the specific configuration. It needs to implement your specific interface and `PlacefulConfigurationConfig`. The `PlacefulConfigurationConfig` interface adds the `getAssociatedInterface()` methods. It returns the "user" interface:

```
SimpleConfigImpl implements SimpleConfig, PlacefulConfigurationConfig {
    private String value;
    public String getValue() {
         return value;
    }
    public void setValue(String value) {
        this.value = value;
    }
    public Class<Simple> getAssociatedInterface() {
        return Simple.class;
    }
}
```

- Create the class for the merge algorithm. It has to implement `PlacefulConfigurationAlgorithm`. Its only method takes a PC and a storage and returns a merged PC. Our Simple configuration will do no merge at all:

```
public class SimpleMergeAlgorithm implements PlacefulConfigurationMergeAlgorithm {
    Simple mergeEntries(Simple pc, PlacefulConfigurationStorage storage) {
        return pc;
    }
}
```

- Finally the "user" interface has to know the implementing and merge class. You add the `@configurationClass` annotation on the "user" interface:

```
@ConfigurationClass(value=SimpleConfigImpl.class, mergeAlgorithm=SimpleMergeAlgorithm.class)
public interface Simple extends PlacefulConfiguration, SimpleConfig, Serializable {};
```

- If you want your PC to be usable by a Directory storage you also need to provide a schema Section 23.4.2, “Directory storage”

In most case you want your merge algorithm to merge all entries from the root to the node of the configuration. For this, you can extend the class AbstractBaseMergeAlgorithm. You then only need to add a method to merge 2 entries.

If your PC is a simple JavaBean, then you can extend BeanMergeAlgorithm . The only things you need to pass is the class of your bean. For each property, the merge algorithm will return the nearest value from the node that is not null.

23.4. Available storage

A storage specifies where the configuration is stored. You can pass specific values to each storage using the properties map (see Section 23.5, “Exemple of extension definition”). The storage gives you also the possibility to define fields that will be indexed. When a user query the storage, only the indexed fields can be queried. At the moment, only String field can be queried and indexed.

It is assumed that an indexed field is a property of the PC and so available via getters. By default, all the "placeful" values (principal, path, repository) are indexed and don't need to be added to the fields list.

23.4.1. In memory storage

The InMemory storage does not use any persistence. You should only use in it in very simple situation such as test. The storageBackend class is: `org.nuxeo.ecm.platform.placeful.configuration.storage.InMemoryPlacefulConfigurationStorage`. You can pass comma-separated values in the fields property. Each value represent a property of the PC. In the backend, an index will be created for each property. You can then use it as a search field in the map passed to manager to find configurations (Section 23.2, “Using Placeful Configuration”).

23.4.2. Directory storage

`org.nuxeo.ecm.platform.placeful.configuration.storage.SQLDirectoryPlacefulConfigurationStorage` implements the storageBackend. You pass it as the class attribute. You also need to add the name of you directory:

```
<storageBackend name="SQL"
    class=
"org.nuxeo.ecm.platform.placeful.configuration.storage.SQLDirectoryPlacefulConfigurationStorage">
  <properties>
    <property name="directoryName">localTheme</property>
  </properties>
</storageBackend>
```

To use directory storage you need to define a directory and the associated schema. The schema only needs to include the `base.xsd` schemaLocation and the `placefulConfiguration.xsd` schemaLocation. You can also add any String element that you want to be indexed. Note that the "placeful" fields general for all configurations are included by default and don't need to be added.

Don't forget to declare the schema and directory to the extension point. It is assumed that each field to be indexed is a property of the PC and can be accessed by getters.

23.5. Exemple of extension definition

```
<component
    name="org.nuxeo.ecm.platform.placeful.configuration.defaultContrib">

  <extension
      target=
"org.nuxeo.ecm.platform.placeful.configuration.service.PlacefulConfigurationService"
      point="storage">                                                          ❶❷

    <storageBackend name="RAM"
        class=
"org.nuxeo.ecm.platform.placeful.configuration.storage.InMemoryPlacefulConfigurationStorage">
      <properties>
        <property name="fields">docId,theme,mode</property>                     ❸
      </properties>
    </storageBackend>

  </extension>

  <extension
      target=
"org.nuxeo.ecm.platform.placeful.configuration.service.PlacefulConfigurationService"
      point="configuration">                                                    ❷
```

```
    <configuration name="TestConfig" storage="RAM"
        class="org.nuxeo.ecm.platform.placeful.configuration.entry.LocalTheme" />

  </extension>

</component>
```

❶ Definition of the storage.
❷ The configuration reference the storage name using the storage attribute.
❸ Storage specific values.

Chapter 24. Monitoring Nuxeo

Nuxeo comes with management by providing a resource publisher, an use case scheduler and a REST XML serializer. Nuxeo makes use of theses services for providing you information about components deployment, directory status, events metric, http and directories session metrics.

These services is to used for integrating nuxeo in a management platform.

24.1. Integrating Nuxeo monitoring in your management system

As an administrator, I want to monitor my Nuxeo server.

Nuxeo exposes monitoring data and behaviors by using the JMX standard. If your management platform does not support JMX, you're able to feed it using the XML serializer provided as a WebEngine module. Nuxeo provided mbeans are published in the "nx" domain.

Nuxeo has identified these kinds of monitoring data : inventory, metric and usecase. Nuxeo types theses objects using the "management naming attribute.

24.1.1. Inventory (nx:*,management=inventory)

The inventory is based onto the component registry. It mainly adapt the registration objects graph (components, extension points, contributions) as an mbean tree. This gives you an access to what and how are the services deployed. Other informations about these services (metrics, status, usecases) will be is to be attached to the inventory. Using these information, you can check that your server is correctly deployed.

24.1.2. Metric (nx:*,metric=*,management=metric)

Inventory is being completed by metrics. It's the accounting information that they can easily provide to users with no costs overhead. As example, directory manager provide core sessions counters. These information are bound into the mbeans tree. Managers defines gauges that polls metric for treesholds or to gather historical data.

24.1.3. Quality (nx:*,usecase=*,management=usecase)

Services are tested periodically by running a typical use case. Results about are published in the service subtree by the management use case scheduler. These informations is to be polled by managers for issuing alerts. We already defined a use case checking that the repository is well working by creating and removing a document.

24.2. Integrating management in nuxeo server

As an administrator, I want to enable nuxeo platform management features.

Management is packaged into three distinct packages : nuxeo-runtime-management, nuxeo-platform-management and nuxeo-webengine-management modules.

24.2.1. nuxeo-runtime-management

That module contains the management logic : the resource publisher and the use cases schedulelr. These services implements only behaviours and have to be contributed by other components for activation.

24.2.2. nuxeo-platform-management

That module contains adapters for nuxeo services such as the runtime inventory or the http session metric. It

also contains the use case scheduler and some typical use cases such as the directory one.

Nuxeo platform management provides you the basic monitoring resources. The following section describes how coding integrating newer monitoring resources suitable for your needs.

24.2.3. nuxeo-webengine-management

That module contains a web engine module for serializing in XML mbeans. You have to use it if your management platform is not JMX aware.

24.3. Contributing management

24.3.1. Publishing

As a developper, I want to publish some informations suitable for monitoring services I'm in charge for.

Here is the typical use cases for : service publishing, resource publishing and mbeans aliasing.

24.3.1.1. Publishing a service

As a developer, I want to publish a service. Usually, services are defined as singleton.

To publish summarized informations about this service, you first have to define a Java interface for and makes your singleton implementing it. Then, you have to contribute to the resource publisher service, naming your singleton. Given the monitoring interface `myPackage.MyServiceMBean` and the implemented class `myPackage.MyService` , you have to define an extension in service definition as follow :

```
          ..
<require>org.nuxeo.runtime.management.ResourcePublisher</require>
<extension point="services"
  target="org.nuxeo.runtime.management.ResourcePublisher">
  <service class="mypackage.MyService"
    ifClass="myPackage.MyServiceMBean" name="myService" />
</extension>
..
```

Note

As you respect a JMX convention, resource publisher is able to guess the interface class you use by its name. So, in that case, it is not mandatory to provide the interface class name.

24.3.1.2. Publishing resources using a factory

As a developer, I want to publish a collection of informations, such as a map of metrics.

The first idea is to expose a getter that returns the map. That kind of information is not well supported on manager side. Monitors and gauges are just able to bind to attribute with a string or numeric type. If the cardinality is acceptable, the resource publisher enables you to register a factory that is to be call backed for publishing your resources. Given the monitoring interface `MyMetricMBean` , the `MyMetricFactory` will publish metrics provided by `MyService`

```
           package mypackage;

import org.nuxeo.runtime.management.ResourceFactory;
import org.nuxeo.runtime.management.ObjectNameFactory;

public class MyMetricFactory implements ResourceFactory {

  public void configure(ResourcePublisherService publisher, ResourceFactoryDescriptor descriptor) {
    this.publisher = publisher;
```

```
    this.service = (MyService)Framework.getLocalService(My.class);
  }

  protected final ComponentName myServiceName = MyService:NAME;

  protected MyService service;
  protected ResourcePublisher publisher;

  public void registerResources() {
    for (String name:myService.doGetMetricNames() {
      doRegisterMetric(name);
    }
  }

  protected void doRegisterMetric(String name) {
    MyMetricContext context = new MyMetricContext(String name, service);
    String shortName = ObjectNameFactorty.formatMetricShortName(name);
    ObjectName qualifiedName = ObjectNameFactory.formatMetricQualifiedName(myServiceName,name);
    publisher.registerResource(shortName, qualifiedName, MyMetricMBean.class, metric);
  }

}
```

```
            package mypackage;


public class MyMetricContext implements MyMetricMBean {

  public void MyMetricContext(String name, MyService service) {
    this.service = service;
    this.name = name;
  }

  protected final String name;
  protected final MyService service;

  int getCount() {
    return service.doGetCount(name);
  }


}
```

```
            package mypackage;


public class MyService implements Service {

   protected String doGetNames() {
     ..
     return ...
   }

   protected int doGetCount(String name) {
     ..
     return ...
   }
}
```

```
                    ..
          <require>org.nuxeo.runtime.management.ResourcePublisher</require>
          <extension point="factories"
            target="org.nuxeo.runtime.management.ResourcePublisher">
            <factory class="mypackage.MyMetricFactory"
              ifClass="myPackage.MyMetricMBean" name="myMetricFactory" />
          </extension>
          ..
```

24.3.2. Providing shortcuts

As an administrator I want to monitor OperatingSystem attributes in my management system using the XML serializer. I will contribute to the resource publisher by defining the following extension point :

```
                  ..
```

```
          <require>org.nuxeo.runtime.management.ResourcePublisher</require>
          <extension point="shortcuts"
            target="org.nuxeo.runtime.management.ResourcePublisher">
            ..
            <shortcut name="operatingSystem" qualifiedName="java.lang:type=OperatingSystem"/>
            ..
          </extension>
          ..
```

24.3.3. Reporting quality

As a developer, I want to provide a feature operational status.

Services quality are to be monitored by scheduling typical use cases. Use cases should throw an exception for indicating error condition. Use cases are registered under the service they belongs to.

```
                ..
          <require>org.nuxeo.ecm.management.usecases.UsecaseScheduler</require>
          <extension point="usecases"
            target="org.nuxeo.runtime.management.ResourcePublisher">
            <usecase name="myUsecase" class="mypackage.MyUsecase"
              serviceClass="mypackage.My" />
          </extension>
          ..
```

```
          package mypackage;

import org.nuxeo.ecm.management.usecases.Usecase;

public class MyUsecase implements Usecase {

   void init(Object service) {
     this.service = (My)service;
   }

   MyService service;

   void runCase(CoreSession session) throws ClientException {
      service.doSomething();
   }

}
```

Part III. Core Services

Chapter 25. Nuxeo Runtime

25.1. Overview

Nuxeo Runtime is the foundation of the Nuxeo infrastructure. It handles deployment and extensibility of components to target platforms. This component allows the whole Nuxeo infrastructure to be easily ported between Java platforms (Java EE, OSGi, etc.) and features an easy plug-in mechanism that any component can use to declare extension points. These extension points can be used by other components to extend the former one.

The Nuxeo Runtime uses the OSGi component model and a set of adapters to deploy POJO components to Java host platforms, such as Eclipse/Equinox, or a Java EE 5 application server such as JBoss or WebLogic. When deployed, Nuxeo Runtime components become actual host platform components. For example on JBoss the component is seen as a MBean, while when deployed on Geronimo it is seen as a GBean and on Eclipse it is seen as a native Eclipse plug-in. In short, Nuxeo Runtime offers a new and seamless way to make your Java EE applications and components extensible (as Eclipse developers are already used to).

Nuxeo Runtime is not specific to the Nuxeo platform, it is a generic deployment and extension system that can be used in any Java or Java EE application.

Forget specific build of your applications for a dedicated project or customer and enjoy "Code once, deploy anywhere" for real!

25.1.1. Main Goals

One of the main requirements of the "Nuxeo Core" component is to be deployable on both the JBoss and Eclipse platforms. To ease development and allow as much code reuse as possible, Nuxeo developed a common component and packaging model that may be deployed and used on both of these platforms without any code change or repackaging. This model, the Nuxeo Runtime, was developed as the foundation for all Nuxeo components. The Nuxeo Runtime is not a standalone framework. It is, in short, a component model running on top of an existing platform that provides a common, platform-independent, model to underlie the components of an application. This architecture allows flexible and true componentization of applications.

In addition to the component model, the Nuxeo Runtime also defines a common model for packaging, the OSGi bundle model. At the lowest level, OSGi bundles are simply regular JARs containing an OSGi manifest file. OSGi technology is becoming more and more popular and is currently used by Eclipse, Geronimo, Glassfish, and Jonas as their runtime framework. Because of the component and packaging choices, applications based on Nuxeo Runtime can run on different platforms without modification and without having to care about the particulars of a deployment platform.

25.1.2. Main Features

The main features provided by Nuxeo Runtime are:

1. Native OSGi support
2. Extensible component model through extension points
3. Adapters to support host platforms (JBoss and Eclipse support is built-in)

25.2. What is OSGi?

OSGi (Open Services Gateway initiative) is an open standards organization founded by Sun Microsystems, IBM, Ericsson and others in March 1999. OSGi defines a modular and complete, Java-based service framework. The deployment units used by this framework are called *bundles*, so we will refer them as *OSGi bundles*.

OSGi bundles are normal Java libraries (JAR files) containing a special *manifest* file (`META-INF/MANIFEST.MF`) describing all aspects related to the bundle, for example: the bundle name, description, bundle dependencies, exported packages, the bundle classpath, the bundle activator and many other OSGi-defined features.

Here is a typical OSGi manifest file:

```
Manifest-Version: 1.0
Bundle-ManifestVersion: 2
Bundle-Name: NxRuntimeEclipseDemo Plug-in
Bundle-SymbolicName: org.nuxeo.runtime.demo.eclipse.Demo; singleton:=true
Bundle-Version: 1.0.0
Bundle-Activator: org.nuxeo.runtime.demo.eclipse.demo.Activator
Bundle-Vendor: Nuxeo
Bundle-Localization: plugin
Require-Bundle: org.eclipse.ui,
  org.eclipse.core.runtime,
  org.nuxeo.runtime.demo.HelloWorld,
  org.nuxeo.runtime
```

The *bundle activator* is a Java object that is called whenever the bundle is started and stopped by the framework. This is the only way available to the application to access the framework functionality.

Besides bundles and bundle management, OSGi provides a *service registry*: an API to register the services provided by a bundle and to lookup services defined by other bundles. Also, OSGi defines a Declarative Services specification that significantly simplifies the service-oriented programming model. Through this model, services can be defined in XML files inside the bundle and automatically deployed by the framework.

For a complete definition of OSGi, see Wikipedia.

A key goals of Nuxeo Runtime is to natively support various implementations of the OSGi framework and to use the OSGi bundle model for packaging and deployment. Another goal is to align the Nuxeo Runtime component model with the OSGi-notion of Declarative Specification.

25.3. OSGi Support

Nuxeo Runtime provides "built-in integration" with any OSGi-compliant framework. This means Nuxeo Runtime-based components can run “as is” on any OSGi-enabled platform. On other platforms like JBoss, an adapter is required. The Nuxeo Runtime eases the creation of such adapters by providing an abstract OSGi adapter that can be customized for any platform.

Note: This does not mean you can transform any platform into a fully OSGi-compliant platform using Nuxeo Runtime adapters. Primarily, this is because the adapter must use the host platform's class-loading and deployment model that may be incompatible with the OSGi specifications. The Nuxeo Runtime's adapters only mimic an OSGi environment, using native host platform features, for applications using Nuxeo Runtime. Many OSGi features are not yet provided by the adapter – but we hope to add more and more features. If you are interested in helping on this, do not hesitate to contact us :-). Currently, one of the most important features that is missing is OSGi service support, but we are working on this and hope to provide it soon.

When running on true OSGi platforms, no adapter is used and thus all OSGi features are available and supplied by the host platform. Nuxeo Runtime components run on such a platform without any alteration.

Currently we provide two built-in adapters:

1. JBoss OSGi adapter – used to deploy OSGi bundles on JBoss AS 4.x
2. Test OSGi adapter – used for JUnit testing and can be used on any simple Java application that is not using a complex class loading or deployment mechanism.

25.3.1. Supported Features

Currently, Nuxeo Runtime adapters can provide the following OSGi features:

1. OSGi Bundle deployment

2. Manifest file loading
3. Classpath processing
4. BundleActivator support (activation and deactivation)
5. BundleActivator notification each time a bundle is started and stopped
6. Fake Bundle and BundleContext implementations that adapt OSGi operations to native operations of the host platform
7. Support for the common operations defined by the OSGi API for BundleActivators
8. Bundle lifecycle and framework support
9. Bundle dependencies (as specified in the manifest)

25.3.2. Unsupported Features

The following OSGi features are not supported (yet):

1. The OSGi service layer
2. The OSGi security layer
3. The OSGi class-loading specifications (the class-loading mechanism of the host platform is used)
4. Some methods of the interfaces Bundle and BundleContext (unimplemented Methods will thrown an `UnsupportedOperationException` exception)

25.3.3. Planned Features

1. Supporting the OSGi service layer
2. Implementing the OSGi declarative services based on the runtime component model

25.4. Component Model

The component model provides a flexible way to define, register and locate components. It was designed in order to reuse the same components on very different platforms like JBoss and Eclipse.

Full support of OSGi declarative service specifications is planned for the medium-term future. In addition to this, Nuxeo components can describe any type of components, not only services.

25.4.1. What are components?

A definition from Wikipedia: “A software component is a system element offering a predefined service and able to communicate with other components”. Components as defined by the Nuxeo model are logical units that may depend on and/or extend one another. The Nuxeo Runtime is responsible for providing a common API to register, locate or extend components and most commonly components are registered using XML descriptor files.

Components can be declared as independent, top-level components - by using a standalone XML file in the bundle- or they can be declared at a finer granularity by programmatically registering sub-components within the bundle. To declare an top-level component, you need to create an XML description file, put it somewhere in the bundle, typically in the OSGI-INF directory, and specify the “Nuxeo-Component” header in the bundle

manifest to load the component at bundle activation. For example, the manifest shown below has a reference to the XML file "helloworld-extension.xml" that declares a component.

```
Manifest-Version: 1.0
Bundle-ManifestVersion: 2
Bundle-Name: HelloWorldExtension Plug-in
Bundle-SymbolicName: org.nuxeo.runtime.demo.HelloWorldExtension
Bundle-Version: 1.0.0
Bundle-Vendor: Nuxeo
Bundle-Localization: plugin
Require-Bundle: org.nuxeo.runtime.demo.HelloWorld
Nuxeo-Component: OSGI-INF/helloworld-extension.xml
```

25.4.2. Main Features

1. *Declarative components through XML descriptors*: XML component descriptors are tightly integrated with OSGi – you can specify which components should be deployed at bundle activation-time by using the custom manifest header “Nuxeo-Component”.

2. *Dependency between components*: components are activated only when all their prerequisites are met. If prerequisites are not met the components will be put in a pending state until their dependencies are completely resolved. Similarly when uninstalling a component, all the components depending on it will be moved to the pending state when the a component is deactivated.

3. *Extensibility through extension points*: each component can let other components extend its functionality by defining a set of extension points. Other components (or the component itself) may then plug extensions (Java code) into one of the declared extension points. This flexible extension mechanism draw inspiration from the Eclipse extension points but is not identical to it.

4. *Life Cycle Events*: component life cycle events are sent by the runtime to any interested component. See Adaptable Components for a common use case.

5. *OSGi integration*: the component model is about to be fully integrated with OSGi and will be soon compliant with the OSGi declarative service model.

6. *Platform Independence*: the component model can be used on any platform. It provides a single API to register and look-up components – the Nuxeo Runtime native API may be used (and in the future, the OSGi service API will be available too).

25.4.3. Planned Features

1. Complete integration with OSGi declarative services specifications

2. Component lookup through JNDI

25.4.4. Adapting Components

The runtime implementation may "adapt" registered components to components of the host platform. (This can often be accomplished using the Nuxeo component life cycle notifications.) The JBoss adapter for the Nuxeo Runtime, for example, is already doing this to adapt runtime components into JBoss "MBean services", and thus Nuxeo "components" are seamlessly integrated into the host platform. Since these components are understood by the host platform, we can leverage existing host platform functionaly, such as MBean service management on JBoss.

25.4.5. Flexible Model

You should not be afraid by the “component model” denomination. The runtime component model does not limit your objects, nor imposing any extra rules on your development practices. You do not need to modify your existing objects to derive from some Nuxeo-supplied base class, nor implement some runtime interfaces in order to plug your objects into the component registry. Your objects that implement a "component" may be of

any kind.

The only requirement is that to be a "component", your object must have a public constructor without arguments (the default constructor) so that it can be instantiated via `newInstance` method on the your object's Class object. The component model makes this a mild restriction, since it supplies other means for object initialization if you need to do complex actions at the time your object is created.

If you want to benefit from the extension points mechanism or to respond to component life cycle events like activation or deactivation, there are two options. You can either implement the `Component` interface if this is convenient for you, or define some methods with a particular signature so your object will be called by the runtime using Java reflection. See below for more details on this.

In conclusion, the component model offers substantial benefits with few requirements; the model primarily provides you the capability to register your components, extend others components, allow others to extend your componentns, and look up components. These functions work the same way on any platform supported by Nuxeo Runtime.

25.4.6. Component Life Cycle

A component has three primary life cycle states:

1. `Registered`: the component registration information was created and inserted into the registry. The component dependencies are not yet processed or resolved, so the component cannot be activated.

2. `ResolvedAll`: dependencies of this component are satisfied. The component can be safely activated.

   Other unresolved components waiting for a resolution that depends on this component are notified and if they have no more dependencies, they will be moved to this state as well.

3. `ActivatedComponent`: Activation occurs as a result of one of three events: immediately upon component becomin resolved, programmatically at the user request, or lazily the first time the component is referred to by another component. The only requirement for activation is that the component must be resolved.

   Currently only the immediate activation mode is supported.

   When an activated component is deactivated it is put back into the resolved state. If a components is to become unregistered, it will first be put in the resolved state, then an unresolved event is fired and the component regresses to the registered sate, and finally it is removed from the registry.

   When a component is activated or deactivated the runtime will call the activate or the deactivate method of the component, if any. Implementing life cycle methods is the programmer's choice, they are not required. These methods can be used to initialize and destroy the component in the given context.

   The activation of a component signifies that the component is available and ready to be used by other components, so that the component must be correctly initialized when it enters this state. The following diagram illustrates the life cycle states, method call, and messages sent for Nuxeo Runtime components.

   a. XXX ADD GRAPHIC HERE

25.4.7. Component Extensibility

A key innovation of the Nuxeo Runtime component model is the extension mechanism that enables components to extend one another.

We will begin with a demonstration of how this mechanism works: Imagine you have a component A, implemented in Java with class ImplA, that manages an action menu for the application. A wants to let other components contribute actions to the menu in an easy and flexible way – for example by using XML files to describe these actions.

To be able to do this, component A should declare an extension point, let's name this point “actions”. Other

components willing to contribute some actions to this "group effort" of a menu should use "actions" to indicate their contribution. (A component may declare any number of extension points and may contribute to any number of extensions to other components' extension points but we are using only point for this example.)

Components may declare extension points and extension contributions using a simple XML syntax like the following:

```
<?xml version="1.0"?>
<component name="A">
  <!-- A 'exposes' an extension point 'actions' that should be supplied with other components' MenuItem objects
  <implementation class="ImplA"/>
  <extension-point name="actions">
    <object class="MenuItem"/>
  </extension-point>
  <!--A is 'contributing' a fragment of XML describing a 'doctype' to the extension point 'documentTypes' of B -
  <extension target="B" point="documentTypes">
    <doctype name="File" extends="Document">
      <schema name="common"/>
      <schema name="file"/>
    </doctype>
  </extension
</component>
```

The content of an extension element is specific to the target extension point, such as the "doctype" in the example above. The extension element content is known only by the extension point. If a contribution to an extension point is made an the xml snippet is not correct from the standpoint of the component exposing the extension point, the result is unpredictable - generally, bad contributions will be ignored and some error will be logged. There is, for now, no mechanism of validating XML extensions like in Eclipse.

The Nuxeo Runtime provides an easy way to map XML extensions to real Java objects through an XML mapping mechanism called XMap. In the example above, the "doctype" can be thought of an object with a field called schema that contains a list of strings. Nuxeo supplies the XMap library to allow the XML fragment to automatically be transformed into a Java object with the fields correctly filled in. When not using XMap, extensions are returned as DOM elements and thus the component should itself perform the parsing of extension contributions.

For details on the XML mapping of XMap, see the XMap documentation and/or the JavaDoc.

XXX TODO ADD GRAPHIC

25.4.7.1. Use Cases

Here is the list of some use cases of the extension mechanism identified in the context of the Nuxeo ECM Platform:

1. to define actions and menus
2. to define content schemas (by importing XSD files)
3. to define views (view ids mapped to JSF/XHTML pages)
4. to define content objects (associate a Content Schema with a class that will provide required methods for the content object)
5. to define permissions (usable in security annotations)
6. to define PageFlows for Seam that, optionally, can extend existing ones
7. to define business processes (in jBPM)
8. to define content transformations (doc -> pdf, doc -> html, odf -> pdf, odf -> html, etc.)
9. to define rules for the rule engine (that can be bound to some objects to run a rule only in a specific folder)

10. scriptable extensions that define scripts binded to interpreters like JavaScript, Groovy, Jython, JRuby, etc
11. to define JMS/Event queues
12. to define event types
13. to define security policies
14. to define Access Control Policies
15. to define NXCore storage backends (JCR, SQL, LDAP, etc.)
16. to define query engines
17. to define indexing engines

25.5. Supported Host Platforms

25.5.1. JBoss Integration

JBOSS permits its extensions to be packaged as SAR files, thus the Nuxeo Runtime provides a SAR package containing the Nuxeo Runtime and the JBoss OSGi adapter. This package is an OSGi bundle that acts as the OSGi system bundle. With the Nuxeo SAR in place inside JBoss, any packaging file format understood by JBoss (.sar, .jar, .ear, .war) or even a raw directory, will be treated as an OSGi bundle if it is found by JBoss and contains a valid OSGi manifest. In order for these bundles to be deployed, you need to have `NXRuntime.sar` already deployed in JBoss.

Besides the OSGi adapter and the auto-registration of components through bundle manifest, the JBoss adapter adds the capability to deploy runtime components as XML files located outside OSGi bundles through the JBoss deployment mechanism. This feature can be useful to register components that provide extensions to other components that can be described by plain XML without any code dependency.

Example of a plain XML component that contributes new document types:

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.core.CoreExtensions">
  <extension target="org.nuxeo.ecm.core.schema.TypeService" point="doctype">
    <doctype name="File" extends="Document">
      <schema name="common"/>
      <schema name="file"/>
    </doctype>
    <doctype name="Folder" extends="Document">
      <schema name="common"/>
      <facet name="Folderish"/>
    </doctype>
    <doctype name="Workspace" extends="Document">
      <schema name="common"/>
      <facet name="Folderish"/>
    </doctype>
    <doctype name="Domain" extends="Document">
      <schema name="common"/>
      <facet name="Folderish"/>
    </doctype>
  </extension>
</component>
```

The `NXRuntime.sar` library offers two JMX services:

1. The adapter service (`nx:service=adapter`)
   a. deploys OSGi bundles and declared components
   b. provides information about deployed bundles and components through the JBoss JMX Console

2. The XML component deployer (`nx:name=bundleDeployer,type=deployer`) that can deploy XML descriptors as OSGi components

25.5.1.1. Installation

Deploy the `NXRuntime.sar` in JBoss, then deploy your OSGi bundles as common JBoss packages in any JBoss-supported format.

That's it,, your bundles are deployed and activated.

25.5.2. Eclipse Integration

For Eclipse, a `NXRuntime.jar` bundle is provided. Since Eclipse is OSGi-compliant, Nuxeo Runtime will not install any adapter (so it is not intervening on the bundle deployment).

When running on OSGi platforms, the main role of the runtime is to register components declared inside OSGi bundles (as seen previously through their manifest).

Because Eclipse is not starting automatically OSGi bundles (it starts them only on demand or on class loading), you need to update Eclipse's `config.ini` and configure it to start Nuxeo Runtime (i.e. org.nuxeo.runtime) when Eclipse starts:

```
osgi.bundles=org.eclipse.equinox.common@2:start, org.eclipse.update.configurator@3:start, org.eclipse.core.runti
```

25.5.2.1. Installation

Update the `config.ini` file as described above, then copy `NXRuntime.jar` inside the Eclipse plugin directory.

Start Eclipse. You're done!

25.6. Using Nuxeo Runtime

25.6.1. Creating Components

Components may be created either from XML descriptor files, or programatically.

In order to register components you always need a runtime context.

25.6.1.1. Runtime Context

A *runtime context* is the context where a component is registered. Contexts are always associated with the bundle containing the component classes. Through the context, a component can access the runtime service and can load classes and retrieve resources from its bundle and other visible bundles. `RuntimeContext` objects depend on the current implementation of the runtime service, which varies based on the deployment environment:

Nuxeo Runtime provides three implementations of the `RuntimeContext` interface:

1. `org.nuxeo.runtime.model.impl.DefaultRuntimeContext`: this is a simple implementation of a context designed to be used outside of an OSGi environment. This context uses the current thread context class loader. It is provided so that simple Java applications like JUnit tests can function properly without needing a full-blown OSGi system.

2. `org.nuxeo.runtime.osgi.OSGiRuntimeContext`: this context can be used on any platform supporting OSGi bundles. This context uses the bundle's ClassLoader to load classes and find resources. This is the context supplied when the Nuxeo Runtime is deployed on a OSGi platform.

3. `org.nuxeo.runtime.jboss.JBossRuntimeContext`: this is a JBoss specific context. This implementation wraps JBoss' DeploymentInfo object and uses the the JBoss infrastructure to load components deployed as standalone XML files.

Once you have a runtime context object, you can start registering components.

25.6.1.2. Creating components from XML descriptor files

To create a component using its XML description, follow these steps:

1. Write the XML description of the component.

   Example of a simple XML descriptor:

```
<?xml version="1.0"?>
<component name="org.nuxeo.runtime.EventService">
  <implementation class="org.nuxeo.runtime.services.event.EventService"/>
  <extension-point name="listeners">
    <object class="org.nuxeo.runtime.services.event.ListenerDescriptor"/>
  </extension-point>
</component>
```

2. Load the XML file and register the component.

   In order to register a component, we always need a runtime context:

```
// retrieve the current bundle
Bundle bundle = ...
// create a context given the current bundle object
RuntimeContext context = new OSGiRuntimeContext(bundle);
// load the component XML file given its location relative to the bundle root
context.deploy("OSGI-INF/MyComponent.xml");
```

Note

The current bundle object is usually retrieved from a `BundleActivator` in the `start(BundleContext context)` method. You can also lookup other bundles by their symbolic names given a `Bundle` object.

The context has several method of deploying (e.g. installing) components. For example, the method used previously (`Context.deploy(String)`) is identical to:

```
// load the component XML file given its location relative to the bundle root
URL url = context.getLocalResource("OSGI-INF/MyComponent.xml");
if (url != null) {
    context.deploy(url);
}
```

25.6.1.3. Automatic deployment of components

The best and recommend way to deploy components is to let the infrastructure deploy them when the bundle is activated.

This can be done by specifying the local paths of the XML description files inside the bundle's `META-INF/MANIFEST.MF` file by using the `Nuxeo-Component` header as in the following example:

```
Manifest-Version: 1.0
Bundle-ManifestVersion: 2
Bundle-Name: HelloWorldExtension Plug-in
Bundle-SymbolicName: org.nuxeo.runtime.demo.HelloWorldExtension
Bundle-Version: 1.0.0
Bundle-Vendor: Nuxeo
Bundle-Localization: plugin
Require-Bundle: org.nuxeo.runtime.demo.HelloWorld
Nuxeo-Component: OSGI-INF/MyComponent.xml, OSGI-INF/MySecondComponent.xml
```

This way, as soon as the bundle is activated, theses two components defined in XML files, will be automatically deployed.

XML component are contained as resource files in that bundle and their path should be specified as relative to the bundle root with no initial slash. The two components' XML files shown in the previous example are located in the `OSGI-INF` directory at the root of the bundle's jar file.

25.6.1.4. Dependencies between components

Sometimes order of deployments between bundles or components is important and needs to be specified.

There are several ways to declare dependencies between them: it involves writing bundle names in the manifest file properties (Require-Bundle, Nuxeo-Require, Nuxeo-RequiredBy, Nuxeo-Component) and component names in contribution files (using the <require></require> tag).

For JBoss developers who do not care about deployment under eclipse: use "Nuxeo-Require" and forget about "Require-Bundle". Include all compile dependencies of your bundle to nuxeo components, and add other requirements if you'd like to make sure your bundle is deployed after another component. In other words, if bundle A lists bundle B on its Nuxeo-Require line, bundle B is deployed first.

For instance, if you'd like to make sure your bundle is deployed after nuxeo-core-api because you have compile dependencies on it and you would like to avoid class loading issue. You should use "Nuxeo-Require: org.nuxeo.ecm.core.api" in your manifest. If you'd like to make sure your bundle is deployed after nuxeo-core, perhaps because you need the default types, you can put "Nuxeo-Require: org.nuxeo.ecm.core".

Warning

Nuxeo-require should not be used between components that could be deployed on separate servers.

If you would like to understand differences between "Nuxeo-Require" and "Require-Bundle":

- Require-Bundle is an osgi notion, so it should be set/tested when deploying components in a real osgi framework like eclipse (Apogée).
- Nuxeo-Require is used only when deploying on jboss, it falls back on Require-Bundle when not set for compatibility reasons.

As a general rule for both properties, you should put all the compile dependencies of your bundle. This will prevent having any class loading issues. Dependencies are transitive: if you depend on nuxeo-core-api and on nuxeo-platform-types-api, you do not need to state nuxeo-core-api as types-api already depends on it.

For Require-Bundle, you may have to add "non-nuxeo bundles" requirements like apache commons-logging to make it work as expected under eclipse. This will obviously fail on jboss as this module in not seen as a bundle in this context: that's why you need to state your requirements in Nuxeo-Require instead.

"Nuxeo-RequiredBy" can be used as "Nuxeo-Require", except that it's the inverse dependency. It's useful when you need to put your bundle before another one. Plus it does not fail when the other bundle is not found, it may be a solution when needing to express dependencies between components that may be deployed on separate servers (although this should never happen if the api module is correctly done).

If your bundle has compile dependencies on jboss-seam, always add "Nuxeo-RequiredBy: org.nuxeo.ecm.war". The war module deployment triggers seam components detection so you need your bundle to be deployed before this is done, otherwise your seam components will not be detected correctly.

"Nuxeo-Component" is not designed to state dependencies, but as contributions are deployed in the given order,

it can be helpful to express dependencies by modifying this order.

The <require></require> tag can be put in components. It is giving fine information about a component requirements. It can be used to control order when overriding another contribution. For instance, if you'd like to make sure your bundle is deployed after nuxeo-core types declarations, because you need to change default core types for instance, you can put <require>org.nuxeo.ecm.core.CoreExtensions<require> at the beginning of your core types contributions file. When contributing to an extension point, you do not need to express dependency to the component declaring the extension point: this requirement is implicit.

As a general note: usually, unless you're overriding an existing configuration, you do not need to order components as services that load them do not make much checks, especially when you're dealing with ordering od=f contributions within the same bundle. For instance, you do not need to make sure the "Folder" core type is deployed after the "dublincore" schema, as long as both are deployed when the repository is opened. Another example: you do not need to make sure the "file" layout is deployed before the "File" ecm type that references it, as long as both are deployed when the page displaying the document layout is loaded.

25.6.1.5. Creating components programmatically

This method of creating components is not recommended since it is internal to Nuxeo Runtime and it depends on the implementation.

Here is an example on how you can use the API to manually register a component. We assume you are running in an OSGi environment and you have a reference to the bundle object containing the component you want to register.

```
// retrieve the current bundle
Bundle bundle = ...
RegistrationInfoImpl ri = new RegistrationInfoImpl();
// create a context associated to the current bundle
ri.context = new OSGiRuntimeContext(bundle);
ri.name = new ComponentName(“my.component”);
// set the class name of the component to register
ri.implementation= “org.nuxeo.runtime.example.MyComponent”;
// register the component
Framework.getRuntime().getComponentManager().register(ri);
```

25.6.2. Using components

25.6.2.1. Responding to life cycle events

When a component is deployed, Nuxeo Runtime will check its dependencies and if all of them are resolved, the component is resolved and activated. (In the future, lazy activation or activation on demand will be supported too). If component dependencies are not satisfied, the component will be put in a pending queue until all of its dependencies are resolved.

When activating a component, the runtime will check if the component defines the `activate` life cycle method and if true, it will call it to get a chance to the component to initialize itself.

The same thing is done when deactivating the component - the runtime will check if the component defines the `deactivate` life cycle method and if true, it will call it to give a chance for the component to dispose itself properly.

There are two ways to define life cycle methods:

1. By implementing the `org.nuxeo.runtime.model.Component interface`

   In this case, a cast to Component interface is performed and the life cycle methods are called.

```
public interface Component extends Extensible {
    public void activate(RuntimeContext context) throws Exception;
    public void deactivate(RuntimeContext context) throws Exception;
}
```

2. By simply declaring a public or protected methods on the component object using the right signature.

In this case the Java reflection mechanism is used to call the methods.

```
public class MyComponent {
    ...
    public void activate(RuntimeContext context) throws Exception {
        ...
    }
    public void deactivate(RuntimeContext context) throws Exception {
        ...
    }
   ...
}
```

25.6.2.2. Looking up components

After a component is activated, it can be retrieved using the Nuxeo Runtime API.

There are several methods to look-up a component:

- Looking up the component by its name:

```
HelloComponent hc = (HelloComponent) Framework.getRuntime().getComponent(
        "org.nuxeo.runtime.demo.HelloComponent");
```

- Looking up the `ComponentInstance` object corresponding to this component. This object is a proxy to the component object:

```
ComponentInstance ci = Framework.getRuntime().getComponentInstance(
        "org.nuxeo.runtime.demo.HelloComponent");
if (ci != null) {
    HelloComponent hc = (HelloComponent) ci.getInstance();
}
```

25.6.2.3. Working with extension points

Now let's take a look at how a component may define an extension point and how other components may use this extension point to contribute extensions.

25.6.2.3.1. Defining an extension point

A component may define any number of extension points. Extension points are identified inside a component by a unique name. We will describe here how to define extensions using the XML descriptor file. Extension points can also by created by hand using the internal API of Nuxeo Runtime but this is no recommended and it is not documented here.

Extension points are specified in the XML component descriptor using the `extension-point` tag. This tag has a required attribute `name` and one or more optional `object` sub-tags.

1. The `name` attribute.

   This should be unique relative to the parent component and is used to identify the extension points inside a component.

2. The `object` sub-tag can be used to define what kind of objects are contributed by XML extensions. These objects will be created from the extension XML fragment by using the XMap engine that maps XML to Java objects through through Java annotations.

   If no `object` sub-tag is specified, the extension will be contributed as a DOM element.

   The `object` tag has a required `class` attribute that specifies the class name of the objects to contribute.

   The object class will be loaded using the context of the bundle that defined the extension point.

Example of a component declaring two extension points:

1. listeners
2. asyncListeners

```
<?xml version="1.0"?>
<component name="org.nuxeo.runtime.EventService">
  <implementation class="org.nuxeo.runtime.services.event.EventService"/>
  <extension-point name="listeners">
    <object class="org.nuxeo.runtime.services.event.ListenerDescriptor"/>
  </extension-point>
  <extension-point name="asyncListeners">
    <object class="org.nuxeo.runtime.services.event.AsyncListenerDescriptor"/>
  </extension-point>
</component>
```

25.6.2.3.2. Contributing an extension

Once a component declaring some extension points has been activated, other components may contribute extensions to that extension point.

To declare an extension, the `extension` tag is used. This tag must contains a `target` and a `point` attribute.

1. `target`

   The target attribute specifies the name of the component providing the extension point

2. `point`

   The point attribute is the extension point name.

The `extension` element may contain arbitrary XML. The actual XML content is recognized only by the extension point to where the extension is contributed. This means you should know the correct format for the extension XML.

For this reason, it is important for components to document their extension points. If the extension point is using XMap to map XML to Java objects, then you can use annotations existing on the contribution object class to know the XML format. These annotations are easy to understand and can be used as well as a documentation for the XML extension format.

If you are familiar with Eclipse extension points, you may wonder why Nuxeo Runtime is not using an XSD schema to define the content of an XML extensions. The reason is simple: because inside our ECM project we need to be able to define any type of XML content - even configuration files from external tools we use like for example a Jackrabbit repository configuration. Defining and maintaining XSD schemas for this kind of extensions would be painful.

Anyway, using XMap to map extensions to real Java objects makes it easy to use extensions.

Here is an example on how a component is declaring some contributions to the previously defined extension points:

```
<?xml version="1.0"?>
<component name="my.component">
  <implementation class=”MyComponent"/>
  <extension target="org.nuxeo.runtime.EventService" point="listeners">
    <listener class="org.nuxeo.runtime.jboss.RepositoryAdapter">
      <topic>repository</topic>
    </listener>
    <listener class="org.nuxeo.runtime.jboss.ServiceAdapter">
      <topic>service</topic>
    </listener>
  </extension>
</component>
```

You can see how the component is declaring an extension to the listeners extension point defined by the component `org.nuxeo.runtime.EventService`

The result of this declaration is that the `EventService` will register two listeners, one listening on events from the “repository” topic, the other on events from the “service” topic.

25.6.2.3.3. Registering contributed extension

Extensions are contributed to the target extension point immediately after the component declaring these extensions is activated. If the target component (the component declaring the extension point) was not yet activated, the contributed extensions are put in a pending queue and they will be contributed as soon as the target component is activated.

A component willing to declare extension points and accept contributed extensions should declare two protected or public methods: `registerExtension` and `unregisterExtension`.

This can be done either by implementing the `Component` interface, or by declaring these methods with their correct signatures on the component object (as we have seen before for the life cycle methods).

These two methods should have the following signature:

```
public interface Extensible {

    public void registerExtension(Extension extension) throws Exception;

    public void unregisterExtension(Extension extension) throws Exception;

}
```

Note that the `Extensible` interface is extended by the `Component` interface.

When an extension is contributed the `registerExtension` method is called with an argument that points to the actual contributed extension as an `Extension` object.

Components should use this method to do something with the extension (usually to register it somewhere).

When the component contributing the extension is deactivated, the Runtime will call the `unregisterExtension` method using the same `Extension` object as a parameter. This gives a chance to the extended component to unregister extensions when they become inactive.

Here is an example of how extensions are registered and unregistered:

```
public class HelloComponent implements Component {
    public final static ComponentName NAME
        = new ComponentName("org.nuxeo.runtime.demo.HelloComponent");

    Collection<HelloMessage> messages = new ArrayList<HelloMessage>();

    public void registerExtension(Extension extension) throws Exception {
        Object[] messages = extension.getContributions();
        for (Object message: messages) {
            HelloMessage msg = (HelloMessage)message;
            this.messages.add(msg);
            System.out.println("Registering message: " + msg.getMessage());
        }
    }

    public void unregisterExtension(Extension extension) throws Exception {
        Object[] messages = extension.getContributions();
        for (Object message: messages) {
            HelloMessage msg = (HelloMessage)message;
            this.messages.remove(msg);
            System.out.println("Un-Registering message: " + msg.getMessage());
        }
    }
...
}
```

You can see how the contributed objects are fetched from the `Extension` object and then registered into a Java Map. These contributions are objects of type `HelloMessage` as defined by the extension point (using the `object` sub-element)

The contributions are also available as a DOM element so you can use this to retrieve contributions in the case you don't use XMap to map XML extensions to Java objects. This DOM element is corresponding to the extension element from the XML component descriptor.

So if you need to retrieve the DOM representation of the extension, you can do:

```
public void registerExtension(Extension extension) throws Exception {
    Element element = extension.getElement();

    // parse yourself the DOM element and extract extension data
    ...
}
```

25.6.3. XML Component Descriptors

In this section we will describe the most important elements composing an XML component descriptor.

You can inspect the XMap annotations on the class `org.nuxeo.runtime.model.impl.RegistrationInfoImpl` to find all elements that may compose an XML component descriptor.

A complete component descriptor may look like this:

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.core.schema.TypeService">
    <implementation class="org.nuxeo.ecm.core.schema.TypeService"/>

    <require>org.nuxeo.ecm.core.api.ServerService</require>
    <require>org.nuxeo.ecm.core.repository.RepositoryService</require>

    <property name="author">Bogdan Stefanescu</property>
    <property name="description">The component description ...</property>

    <extension-point name="doctype">
        <object class="org.nuxeo.ecm.core.schema.DocumentTypeDescriptor"/>
    </extension-point>
    <extension-point name="schema">
        <object class="org.nuxeo.ecm.core.schema.SchemaBindingDescriptor"/>
    </extension-point>

    <extension target="org.nuxeo.ecm.core.api.ServerService"
               point="clientFactory">
        <factory class="org.nuxeo.ecm.core.api.impl.LocalClientFactory"/>
    </extension>
    <extension target="org.nuxeo.ecm.core.schema.TypeService" point="schema">
        <schema name="common" src="schema/common.xsd"/>
        <schema name="core-types" src="schema/core-types.xsd"/>
        <schema name="file" src="schema/file.xsd"/>
    </extension>
</component>
```

Each component is defined inside its own file. As you can see the root element is `component`. This element has a required attribute `name`. Apart this, all other sub-elements are optional.

Here is a list with all supported sub-elements:

1. `implementation`

   This element is used to specify the component implementation class. The element is not required since one may define plain XML components only for contributing some extensions to other components. We will refer to these components as extension components.

2. `require`

   This element can be used to specify dependencies on other components. The component will be resolved and activated only after all these dependency are resolved.

3. `property`

   This element can be used to define random properties that will be available later to the component when it will be created.

4. `extension-point`

   This element is used to declare extension points. A component may declare any number of extension points.

   See more details on this in Working with extension points section.

5. `extension`

   This element can be used to declare extensions to other components (or to the current component itself).

   See more details on this in Working with extension points section.

25.7. Integration tests for Nuxeo Runtime applications

While it's obviously a good thing to unit-test one's code, it's usually not enough for a module designed to be ran as part of a Nuxeo Runtime application. It's indeed likely in this situation that the module will depend on services provided by other modules, and even maybe on their default configuration. It's even not uncommon that a project-specific module consists almost exclusively of calls to generic services provided by the base software platform.

25.7.1. The NXRuntimeTestCase base class

`org.nuxeo.runtime.test.NXRuntimeTestCase` is a base class to write JUnit tests for Nuxeo Runtime applications. It sets up the Nuxeo Runtime environment (in the `setUp` method) and provides methods to control bundle and resources loading. It is designed to behave in the same manner in Maven and Eclipse situations. Therefore resources must be accessed in a way that does not depend on the actual ordering of classpath.

25.7.1.1. Loading a bundle

To load a whole OSGI bundle, use the `deployBundle` method, whose parameter is the bundle symbolic name, as specified in its manifest.

25.7.1.2. Loading one contribution from a bundle

Loading a whole bundle can be too heavy, or bring unwanted default configurations. Therefore, the `deployContrib` method is provided to load just a resource (service definition, extension point contribution, etc.) from a given bundle. It takes two arguments: the bundle symbolic name, and the path to the contrib from the top of bundle.

25.7.1.3. Loading a test resource

For resources from the test packages, just make an OSGI bundle of the test package, which can be done by creating the `META-INF/MANIFEST.MF` at the top of the target jar, and use `deployContrib` as above.

25.7.1.4. Sample usage

The following is an excerpt from `org.nuxeo.project.sample`:

```
public class TestBookTitleService extends NXRuntimeTestCase {

    private BookTitleService service;

    private static final String OSGI_BUNDLE_NAME = "org.nuxeo.project.sample";

    private static final String OSGI_TEST_BUNDLE = "org.nuxeo.project.sample.tests";

    public void setUp() throws Exception {
        super.setUp();
        // deployment of the whole nuxeo-project sample bundle
        deployBundle(OSGI_BUNDLE_NAME);

        service = Framework.getService(BookTitleService.class);
```

```
    }

    public void testServiceContribution() throws Exception {
        // Lookup is ensured simply by making the 'test' sub-hierarchy a
        // bundle of its own, with a MANIFEST file
        deployContrib(OSGI_TEST_BUNDLE, "sample-booktitle-test.xml");
        assertEquals("FOOBAR Test", service.correctTitle("foobar"));
    }

}
```

25.7.2. Frequent patterns

While working on an integration test, it is always worthwhile to prescribe clearly what is to be tested: either API calls to services provided by other modules, consistency with configuration provided by other modules, default configuration for the module being tested. This is especially important for the non-regression aspect of the testing and the cost of maintaining the tests.

The use-cases discussed below require basic knowledge of the Nuxeo ECM framework. Implicitely, it is somehow assumed that the tested code has to interact with one service. In case of multiple target services, one would have to choose a pattern for each of them.

25.7.2.1. Integration test against base services with testing configuration

We want to check that the API calls from the tested component to other components have the desired effect, but we don't want to rewrite the tests each time the default configuration of the other components change. Typically, this means that we need to deploy the xml contributions that define the services we need, together with the minimal configuration to tie it up together.

Example: the search service is able to configure its indexes automatically from the schemas and core types declaration. We don't want to have to update tests if someone changes the default config that ships with nuxeo-core. Ideally, this test should use `deployBundle` to set up core services, test repository, etc. and then work on dedicated schemas and core types that are loaded by `deployContrib`.

25.7.2.2. Integration test against base services and their default configuration

One can imagine here a core event listener that uses a given schema, a component that needs access to the search service to manipulate some specific documents...

In this case, we need to load the base service and its configuration exactly as they are in the real application and we do want the test to catch errors that are due to a change in said configuration. In this pattern, we'd use `deployBundle` all over the place.

It's likely however that one does not want the test to rely on the default (if any) configuration of the module being tested. If the tested component doesn't carry its configuration but still needs to be deployed within Nuxeo Runtime, `deployBundle` can be used on itself, and then `deployContrib` for the test configuration, after the test package has been upgraded to an OSGI bundle.

Variant: testing of a component and the configuration that comes along. Just think of your tested module as a "base service."

25.7.2.3. Reusing test resources from another component

This is usually neither possible nor recommended. Such situations do appear in the Nuxeo code base, and the proper solution is to provide the wished resources or classes from a `package`, precisely like `org.nuxeo.ecm.runtime.test` does.

25.8. Detailed Architecture

25.9. References

1. Nuxeo.org website: http://www.nuxeo.org/

2. OSGi website: http://www.osgi.org/

3. JSR 277, 291 and OSGi, Oh My! - OSGi and Java Modularity, presented by Richard S. Hall at ApacheCon Europe 2006: http://docs.safehaus.org/download/attachments/2995/osgi-apachecon-20060628.pdf

Chapter 26. Nuxeo Core Documentation

26.1. TODO: BS

Start from the previous documentation (OOo) and update it. :-)

This chapter targets developers that would like to use directly Nuxeo Core.

26.2. Overview

Nuxeo Core is the foundation of the Nuxeo ECM project. It defines and provides all the basic services and functionalities needed to build a complete ECM platform and applications:

1. a repository model,
2. schema and document type management,
3. a query service,
4. a security model,
5. a document life cycle service,
6. a flexible core event service.

Like every Nuxeo ECM component, Nuxeo Core is running on top of Nuxeo Runtime which defines an OSGi-compatible component model.

26.2.1. Main goals

The main goals of Nuxeo Core are:

1. to provide the common services needed to build a state-of-the-art ECM platform,
2. to be accessible both remotely or locally (i.e: to provide a common API accessible both from a remote JVM or directly on the local one),
3. to be deployable anywhere without any modification (through Nuxeo Runtime): in a Java EE application server like JBoss, or embedded in a desktop application like an RCP Eclipse application,
4. to be extensible and flexible; this is inherited from Nuxeo Runtime which provides an extensible component model.

26.2.2. Nuxeo Core Components

Nuxeo Core is composed by the following components:

1. NXCore: core ECM model, services and default implementation
2. NXCoreAPI: defines a client API for NXCore
3. NXCoreFacade: Java EE facade
4. NXJCRConnector: JCR storage backend that leverages jackrabbit this is the default NXCore storage backend

All these components are running on top of Nuxeo Runtime.

26.3. Nuxeo Core Architecture

The Nuxeo Core top level components are all roughly following the same style of development, which is structured in three layers:

1. model layer,
2. implementation layer,
3. facade layer.

There are also a number of services used by top-level components to provide them with common functionalities like the schema service, query service, life cycle or security. These services are simple and cannot operate on their own – they need a context to operate on. These services are exposed through top-level components and may not follow the layering presented below.

26.3.1. Model Layer (or Internal API)

Top level components provide a model (an API) that is internal to the Core – this means that they are not directly accessible from remote JVMs and should not be directly used by clients.

The model provides a generic API that defines the concepts used by the service and that may have several implementations (using different storage backends for example).

Usually this API cannot be accessed remotely since implementations may use local resources that cannot be sent over the network.

For example, the Repository model defines objects like Document, Property, Session, etc. The JCR-based implementation for the Repository model is directly wrapping JCR (Jackrabbit) nodes that cannot be detached from the local JVM and sent over the network.

26.3.2. Implementation Layer

Each service may have one or more implementations for their model. For example, the Repository service may have several implementation for the model it defines – this could be a JCR-based implementation, an SQL-based one, or something else. The same goes for the Directory service, it defines a model that could have an SQL-based implementation or an LDAP-based one.

Implementations may use very specific resources and configuration, and are hidden by the common model defined by the service. This means that implementation-specific objects or APIs are never used directly by other Core components, they are only accessed by the implementation of the internal API.

26.3.3. Facade Layer (or Public API)

On top of their model, components usually define a facade layer that enables external clients to remotely access service implementations.

This layer is also named the Public API because it defines the API exposed to clients. Any client, local or remote, must use the public API of the component, and must not make calls to the internal API.

The main requirement of the public API is to use only serializable objects that can be sent over the network and reconstructed on the client machine.

26.3.4. Deployment

The architecture presented above makes it possible to access the Core services when the Core is running inside the same JVM as the client application (e.g., when embedded in a desktop application) but also when it is on a remote JVM (e.g., deployed as a module inside an application server). In both cases the Core services are accessed in the same way – through the public API.

26.3.4.1. Local Access

26.3.4.2. Remote Access

26.3.5. Client Session

Usually a client opens a session on a Core service through the facade and can then send requests to the Core service until it closes the session.

While a client is connected to a Core service, the latter should track the client session and restore its state (if any) at each client request. When the session is closed by the client, the Core service releases any resource held by that session.

Any data passed between the client and the Core service is serializable and so it can safely be sent over the network. In this way a client can operate identically when running on the same JVM or when running on a remote one.

26.4. The Repository Model

The repository model is the main functionality provided by the Core; it represents the very raison d'être of the Core. Most of the other Core services were written as auxiliary components to perform specific needs of the repository model or to enrich it.

The repository model, as its name suggest, is describing a software component for managing repositories of documents. Repositories store documents in a tree-like structure that enables grouping documents inside folders in an hierarchic manner.

Besides storage, the repository provides functionalities like:

1. document versioning,
2. security management,
3. document life cycle,
4. annotations,
5. SQL-like query.

26.4.1. Document and Schemas

Documents are structured objects described by a set of *properties*. These properties may be used to store document meta-data (e.g., creation date, author, state, etc.) or the document data itself (e.g., binary or text files, attachments, etc.).

The properties that a document may have and their types and constraints are defined through several *schemas*.

The repository model natively supports XML Schemas to define document schemas.

A schema is therefore the way the structure and contents of a document is defined. Through schemas you can usually specify things like:

1. what properties are allowed,
2. the type of each property,
3. the default value of the property, if any,
4. the restrictions on the property values, if any,
5. whether a property is mandatory or not.

In order to create and use documents, you first need to define their structure. For this, you have to define a document type. Then you can create instances of documents of this type.

In some ways, document types and schemas are similar to Java classes and interfaces. A document type may implement some schemas in the same way that Java classes implement interfaces, and a document type can extend another document type in the same way that a Java class can extend another class.

Document Types define one or more schemas that the document structure must satisfy and some other extra properties like facets which will be discussed later.

In conclusion, the unit of work in a repository is the document. To create a new document, you must specify the document type and a path. You can either use existing document types or register new types as we will see in the Extension Points section.

For more information on document types and schemas see the section XXX.

26.4.2. Document Facets

A *Facet* is a behavioral property of a document. As schemas define the document structure and content, facets are used to describe behaviors or capabilities of a document.

For now, facets are simple strings attached to a document type to specify a capability for documents of that type. In the future, facets may evolve to more complex structures, for example to dynamically provide interfaces to manipulate documents according to a capability they offer.

Currently the Core defines two facets:

1. `Folderish`: adds folder capability to a document, so that it can have zero or more children documents,
2. `Versionable`: adds versioning capabilities to a document.

26.4.3. Document Annotations

As we've seen, a document structure and content is strictly defined by the schemas its type implements. But there are many situations where some application-specific data need to be dynamically attached to the document and retrieved later without having to modify the document schemas.

This is very useful for repository extensions that needs to store placeful (i.e., location-sensitive) information on a document – information that cannot be specified by any document schema since its type is not necessarily known in advance.

Annotations are not required to be stored through the same data storage as the document itself. For example one may choose to store document in a Jackrabbit-based repository and to store annotations in a dedicated SQL database.

These annotations usually keep some internal state or data about the document. For example, a tool that may use annotations is the workflow service.

26.4.4. Document Access Control

Usually, manipulating documents requires a set of privileges to be granted to the current user. Privileges given to a user over a document are very dependent on the current context and on the document itself.

Usually, privileges depend on:

1. the document location (i.e., privileges are placeful),
2. the access rules defined on document parents in the hierarchy,
3. the document state,
4. and generally on any rule that was defined over a particular location on the document parents.

Privileges are a standard example of extra information that need to be stored on the document in a placeful manner, so it may be a perfect candidate for the annotation service.

But since privileges are very dynamic and may require expensive computations on every document that is accessed, a separate Security Service exists to manage the storage as it sees fit - and not necessarily through annotations on the document. This is more efficient from a performance point of view.

In the following subsections, we will see what type of information is stored on the document to enforce security and how security checks are done. To ease comprehension of security concepts and evaluation, we will begin the presentation from the smallest unit of security information to the largest one that is stored at the document level.

26.4.4.1. Access Control Entry (ACE)

This is the smallest unit specifying a security rule. It is a very simple object containing three fields:

1. principal: an authenticated entity. For example the user that opened the session on the repository is a principal – but a principal may also be a group of users.
2. permission: the kind of action that may be granted or denied for a principal. This may also be a group of permissions. This corresponds to the Java concept of privilege.
3. granting: specifies whether the given permission is granted or denied to the given principal.

Examples:

1. DENY, John, Read: an access entry that specifies that the reading is denied for the principal John.
2. GRANT, Developers, Drink: an access entry that specifies that drinking is granted for any principal from the developer group.

26.4.4.2. Access Control List (ACL)

An ACL is an ordered list of ACEs. This means it represents a set of access rules. Why ordered? Because usually when evaluating access rules the order is important. This is because evaluation stops on the first DENY or GRANT rule that match the criteria check.

Here is a simple example showing how ordering may influence the security checks. Suppose that we have a principal John that belongs to the Readers group, and an ACL that contains the following two ACEs:

1. DENY, John, Read
2. GRANT, Readers, Read

Suppose we want to check whether principal John is granted reading. Every entry in the ACL is checked (in the order they were defined) and if an entry matches the security check the evaluation stops. Using the example above, John will be denied reading even if it is a member of the Readers group. But if you swap the order of ACEs in the ACL, John will be granted reading.

26.4.4.3. Access Control Policy (ACP)

An ACP is an ordered list of ACLs. Each ACL stored in the ACP is uniquely identified by a name. The ordering is important when security is checked – ACLs at the beginning of the list will be checked first.

The ACP is the object containing the security information that is attached to a document.

Note that ACLs are inherited so that a document will inherit any defined ACLs from its parents in the hierarchy. Inherited ACLs are evaluated after evaluating the local ACLs and from the nearest parent to the remotely related parent.

You may wonder why an ACP is containing several ACLs? And what about ACL names? In a typical situation where security information may only be changed by an administrator through a user interface, a single ACL is enough.

But a complex application may have complex rules to set privileges according to the current document state or context. This is the case for a workflow engine which may decide to revoke or grant privileges depending on the document state or the context.

This means that access rules are changed not only by administrators but also by services like the workflow. To avoid collisions, every tool that needs to change access rules may use its own (named) ACL for setting these rules. If the workflow service considers that its rules are more important than the ones explicitly set by the administrator, it simply places its ACL before the one reserved for the administrator so that it will be evaluated first.

Currently there are two predefined ACLs:

1. local: the local ACL

   The local ACL is the only ACL an administrator may explicitly change through the User Interface.

2. inherited: the inherited ACL

   This ACL is computed each time a security check is performed (unless caching is used). The inherited ACL is the ACL obtained by merging all existing ACLs on the document's hierarchy. This ACL is appended to the ACL list, so it will be evaluated last.

So from a simple security unit like the ACE we end up with a sophisticated structure like inheritable ACPs.

These use cases are not artificial, they are real use cases that a mature ECM product should satisfy.

26.4.4.4. Evaluating Privileges

The evaluation mechanism has been described above. Here is an example of how an evaluation is done.

Let's say the principal John is trying to edit the document D. Editing a document requires the Write permission. Suppose the document D has the path /A/B/C/D – it is a child of the document C which is a child of the document B which is the child of the document A.

To decide if the principal John can edit this document the following steps are taken:

1. The merged ACP for the document D is computed. This ACP is the local ACP set on the document D merged with all parent ACPs. ACLs imported from the parents are appended to the local ACLs so that they will be evaluated at last.

2. Each ACL is evaluated in respect to the order defined by the ACP.

3. Each ACE is evaluated in respect to the order defined by the ACL.

4. If an ACE match a security rule regarding the principal John (or a group which it belongs) and the permission Write (or a permission group from which Write belongs) then the evaluation ends and the access right of the matching ACE is returned

5. If no matching ACE is found then the privilege is denied.

26.4.5. Life Cycle

Within organizations, documents are often regulated. At a given time, a document has a state or is within a phase. The way the document transitions in compliance with regulations from one state to another (or from one phase to another) is in most of the cases defined and managed by business processes or workflows.

Nuxeo Core itself doesn't embed a workflow engine, or still a BPM engine, as such. It only provides a generic way to define document life cycles, the way the document properties related the life cycle are stored and a way to specify which document types follow which life cycles at deployment time.

Thus, the workflow engine that will be deployed along with Nuxeo Core will leverage the API exposed by Nuxeo Core to set the life cycle properties.

The APIs defined in Nuxeo Core regarding life cycle are highly inspired from the JSR-283 specifications that are still in a draft state at the time of writing this document.

Another advantage of such a design is the fact that the life cycle state of a document will be independent of the application (i.e.: workflow variables) and will be embedded within the document itself at storage time, and thus will be exported along with the document properties.

Nuxeo provides a BPM engine that knows how to leverages the Nuxeo Core life cycle API. See http://www.nuxeo.org.

26.4.5.1. Example of document life cycle

Here is a typical lifecycle schema example: XXX TODO.

26.4.5.2. Life cycle definition

Nuxeo Core allows one to define life cycle using extension points. (See the Nuxeo Runtime documentation for more information about extension points.). You will find at the end of this document the complete list of extension points defined by the core, you will find an example of life cycle definition there using the life cycle definition extension point.

The life cycle model defined by Nuxeo Core is simple stateful, or state-transition engine. Including the following elements:

1. Life cycle definition

2. Life cycle state definition

3. Life cycle state transition definition

Again, here, no policy regarding transitions are specified. The workflow or BPM engine will deal with this. Here are the reasons:

1. It gives more flexibility regarding the policy that needs to be applied on the documents by letting dedicated BPM engines deal with that. Thus this is possible to choose which workflow engine to use for your application. (see NXWorkflow)

2. Current JCR specifications doesn't include a default policy model regarding life cycle so it appears logical to not include this ourself at this layer of the architecture

3. It simplifies the model

This is important to note that the life cycle definition is fully independent from the document types themselves which allows the reuse if life cycle for different document types.

26.4.5.3. Life Cycle Manager

The life cycle manager is responsible of the storage of the life cycle related properties. One could think of storing the life cycle property within the JCR, which is the default implementation provided by NXJCRConnector, or still one could think about storing it in a separated RDBMS apart from the content storage.

Because of this, Nuxeo provides an abstraction for this storage allowing one to define a life cycle manager per life cycle definition.

Let's take a look at the life cycle manager exposed by Nuxeo Core:

You can see that the interface is fairly simple. It basically, only specifies how to store and retrieve the state and the life cycle policy of a given document.

For an example of JCR storage see the JCRLifeCycleManager definition :

http://fisheye.nuxeo.org/browse/~raw,r=4233/nuxeo/ECMPlatform/NXJCRConnector/trunk/src/org/nuxeo/ecm/core/jcr/JCL

Note this is how the JSR-283 current specifications specifies the life cycle storage repository side.

You can register your own life cycle managers using the lifecyclemanager extension point defined on the Nuxeo Core side. See the extension points chapter of this document for an example.

26.4.5.4. Document types to life cycles mapping definition

When your life cycle definitions are defined and you did specify the life cycle managers which will take care of the storage you will then need to specify associations in between document types and life cycle.

To achieve this, Nuxeo Core defines an extension point allowing one to specify, independently from the document type definition, such an associations. Please, check the example at the end if this document.

26.4.5.5. Core life cycle service

Nuxeo Core defines a dedicated life cycle service that is used by the Nuxeo Core internals. This service is not exposed at the facade layer because we don't need it there. This service is manipulating directly the repository document themselves (not references and thus is not suitable for remoting purpose).

Actually, the document model itself has been extended so that you can directly invoke this service through the document session itself at facade layer. See next chapter for an overview of the API.

This service is defined under this namespace `org.nuxeo.ecm.core.lifecycle.LifeCycleService`.

26.4.5.6. The life cycle document API and the exposure at the facade layer

The document model exposes a life cycle related API. You can take advantage of this API from the document itself if you are working at core level. Here is the API:

26.4.5.7. Core events and listeners

Nuxeo Core defines a service dedicated to core events. This service is only responsible of core events and allows third party code to register listeners that will get notified when events occur (and that can take specific actions themselves).

This service doesn't take advantage of event service such as JMS or still the NXRuntime event service at this

level because it needs ti be really fast at event processing to not decrease the repository performances for instance.

By using event listener extensions, you can hook up and bridge on another synchronous or asynchronous messaging systems. Let's take some examples.

1. Nuxeo Core defines a bridge to Nuxeo Runtime forwarding events on the NXRuntime event service in an asynchronous way. It defines like that a local event loop shared by all components running on top of NXRuntime.

2. The NXEvents component, not part of the Nuxeo Core, registers a JMS listener bridging Nuxeo Core events to a dedicated JMS topic. It allows message driven beans in the Nuxeo Enterprise Platform to get the Nuxeo Core events. (for instance NXAudit)

You could define whatever listeners you need to forward the Nuxeo Core events on an external messaging system. See the end of this document for an example of such a registration.

26.4.6. Query Engine

The query engine is designed to provide an SQL-like language, called NXQL, to perform document and directory queries.

NXQL offers standard SQL functionality to search records, but can also take advantage of the hierarchical nature of the content repository to provide path-based searches. NXQL is used as the uniform query syntax to access several kinds of repositories. The query engine itself must process and optimize the query, and dispatch it to the different backends and tables that are referenced in the query.

Updates or creation statements are not covered and must be performed through the repository API.

For more information about the query engine, refer to the document about NXQL.

26.4.7. The Public API

As we've seen the internal repository model is not remotely accessible. Because the Nuxeo Core deployment model requires supporting both local and remote clients, the APIs are separated between an internal API and a Public API, designed to fulfill the deployment needs. Any client should use the public API to connect to a Nuxeo Repository.

This public API has only one limitation: any object transferred between the client and the core must be serializable. This way it can be sent over the network and restored on the client side.

So the public API is in fact is a serializable view of the repository model. This has a performance drawback compared to the internal API since it should transform any model object like a Document into a serializable form, but has the benefit of being totally independent from the JVM where the Core runs.

The main interfaces composing the public API are the:

1. `DocumentModel`: the serializable view of a Document.

2. `DataModel`: the serializable view of a document subpart described by a schema.

3. `CoreSession`: a session to the Core repository.

4. `CoreInstance`: the gateway to the Core. It uses session factories to create new sessions (connections) to the Core.

26.4.7.1. DocumentModel

The document model is a data object that completely describes a document. You can see it as a serializable

view of document.

Apart from being a data object, this object also provides some logic. For example a document model is able to lazily load data from the storage if not already loaded or it may check permissions for a given user on the document it represent.

The data contained by document model is grouped in `DataModel` objects.

For each document schema, there is a `DataModel` that contains concrete data as specified by the corresponding schema. You can see a `DataModel` as a data object described by a schema (i.e. a schema instance).

A document contains also data that is not defined by schemas like its internal ID, its name, its parent etc. Thus, apart from these data models there is some information stored as members on the document model like the document ID, the document name, a reference to the parent document, the ACP information (used for security checks), the session ID etc.

Also the data model contains the list of facets that the document type defines.

One of the most important ability of the document model is to lazily load data the first time the data is required from the client. This feature is important because the document may contain a lot of schemas and fields and it will be a performance problem to load all these data from the storage each time a document model is created.

Usually, the client application is using only few `DataModel` fields like the `Tile`, `Description`, `CreationDate`, etc. These are the fields commonly displayed by a tree – like explorer of the repository.

When the client is displaying or editing the document properties – then the document model will load missing data models.

To achieve this, there are schemas or fields that are declared to be lazily loaded. When creating a document model from a document, only the non-lazy schemas and fields are fetched from the storage. For example, a blob field will be always lazy.

26.4.7.2. DataModel

As detailed above, the data model is an object containing the concrete data for a document schema.

Each data model is described by the schema name and the map of fields. The data model contains no logic, it is a pure data object.

Apart from the fields map, the data model contains information about dirty fields (fields that have been modified by the client), so that when saving changes to the repository only modified fields are saved.

26.4.7.3. CoreSession

The `CoreSession` is a session to the Nuxeo Core. The session is opened and closed by a client and gives the client the possibility to interact with the Core.

The Core a session connects to can be located in a separate JVM or in the current one. To create remote or local sessions, you need to use a specific `CoreSessionFactory` object. These objects are usually specified using extension points but you can also use them programatically.

After creating a session, you can begin to retrieve and modify documents through the API exposed by the `CoreSession` interface.

Example of creating and using a session:

26.4.7.4. CoreInstance

This is the gateway to a Core instance. As mentioned above, the Core may be located in a remote JVM. The `CoreInstance` uses `CoreSessionFactory` objects (declared through extension points) to connect to a Core instance and to create a session.

26.4.8. Integration with Applications Servers

The repository is plugged into an application server using the a resource adapter as specified by the J2EE Connector Architecture (JCA).

The resource adapter is write over the repository model so it is not dependent on the repository implementation (like for example JackRabbit).

Currently the resource adapter has been tested only on JBoss AS

The resource adapter enables the repository to take part on transactions managed by the application server.

26.5. Extension Points

This section aims to cover all existing extension points defined by core components and to give some examples of creating new extensions.

26.5.1. Session Factories

Declaring component: `org.nuxeo.ecm.core.api.CoreService`

Extension point name: `sessionFactory`

This extension points is for registering new session factories. Session factories are used to create new Core Sessions.

Currently two session factories are provided:

1. a local session factory – that create sessions to a local Core (that is running in the same JVM as the client)
2. a remote session factory – that create sessions to a remote Core (running in a remote Application Server)

Example 26.1. Example Title XXX

26.5.2. LifeCycle Managers

Declaring component: `org.nuxeo.ecm.core.lifecycle.LifeCycleService`

Extension point name: `lifecyclemanager`

This extension points is for registering new life cycle managers. A life cycle manager is responsible for managing and storing document life cycle information.

Example 26.2. Example Title XXX

Chapter 27. Nuxeo Core Import / Export API

The import / export service is providing an API to export a set of documents from the repository in an XML format and then re-importing them back.

The service can also be used to create in batch document trees from valid import archives or to provide a simple solution of creating and retrieving repository data. This could be used for example to expose repository data through REST or raw HTTP requests.

Export and import mechanism is extensible so that you can easily create you custom format for exported data. The default format provided by Nuxeo EP is described below.

The import / export module is part of the `nuxeo-core-api` bundle and it is located under the `org.nuxeo.ecm.core.api.io` package.

27.1. Export Format

A document will be exported as a directory using as name the document node name and containing a `document.xml` file which hold the document metadata and properties as defined by document schemas. Document blobs if any are by default exported as separate files inside the document directory. There is also an option to export blobs inlined as Base64 encoded data inside the `document.xml`.

When exporting trees document children are put as subdirectories inside the document parent directory.

Optionally each service in nuxeo that store persistent data related to documents like the workflow, relation or annotation services may also export their own data inside the document folder as XML files.

A document tree will be exported as directory tree. Here is an example of an export tree containing relations information for a workspace named `workspace1`:

```
+ workspace1
    + document.xml
    + relations.xml

    + doc1
      + document.xml
      + relations.xml

    + doc2
      + document.xml
      + relations.xml
      + file1.blob

    + doc3
      + document.xml
```

27.1.1. document.xml format

Here is an XML that correspond to a document containing a blob. The blob is exported as a separate file:

```
<?xml version="1.0" encoding="UTF-8"?>

<document repository="default" id="633cf240-0c03-4326-8b3b-0960cf1a4d80">
  <system>
    <type>File</type>
    <path>/default-domain/workspaces/ws/test</path>
    <lifecycle-state>project</lifecycle-state>
    <lifecycle-policy>default</lifecycle-policy>
    <access-control>
      <acl name="inherited">
        <entry principal="administrators" permission="Everything" grant="true"/>
        <entry principal="members" permission="Read" grant="true"/>
        <entry principal="members" permission="Version" grant="true"/>
        <entry principal="Administrator" permission="Everything" grant="true"/>
      </acl>
    </access-control>
  </system>
```

```
  <schema xmlns="http://www.nuxeo.org/ecm/schemas/files/" name="files">
    <files/>
  </schema>
  <schema xmlns:dc="http://www.nuxeo.org/ecm/schemas/dublincore/" name="dublincore">
    <dc:valid/>
    <dc:issued/>
    <dc:coverage></dc:coverage>
    <dc:title>test</dc:title>
    <dc:modified>Fri Sep 21 20:49:26 CEST 2007</dc:modified>
    <dc:creator>Administrator</dc:creator>
    <dc:subjects/>
    <dc:expired/>
    <dc:language></dc:language>
    <dc:rights>test</dc:rights>
    <dc:contributors>
      <item>Administrator</item>
    </dc:contributors>
    <dc:created>Fri Sep 21 20:48:53 CEST 2007</dc:created>
    <dc:source></dc:source>
    <dc:description/>
    <dc:format></dc:format>
  </schema>
  <schema xmlns="http://www.nuxeo.org/ecm/schemas/file/" name="file">
    <content>
      <encoding></encoding>
      <mime-type>application/octet-stream</mime-type>
      <data>cd1f161f.blob</data>
    </content>
    <filename>error.txt</filename>
  </schema>
  <schema xmlns="http://project.nuxeo.com/geide/schemas/uid/" name="uid">
    <minor_version>0</minor_version>
    <uid/>
    <major_version>1</major_version>
  </schema>
  <schema xmlns="http://www.nuxeo.org/ecm/schemas/common/" name="common">
    <icon-expanded/>
    <icon/>
    <size/>
  </schema>
</document>
```

You can see that the generated document is containing one [system] section and one or more [schema] sections. The system section contains all system (internal) document properties like document type, path, lifecycle state and access control configuration. For each schema defined by the document type there is a schema entry which contains the document properties belonging to that schema. The XSD schema that correspond to that schema can be used to validate the content of the schema section. Anyway this is true only in the case of inlined blobs. By default, for performance reasons, the blobs are put outside the XML file in their own file.

So instead of encoding the blob in the XML file a reference to an external file is preserved: cd1f161f.blob

Here is how the same blob will be serialized when inlining blobs (an option of the repository reader):

```
  <schema xmlns="http://www.nuxeo.org/ecm/schemas/file/" name="file">
    <content>
      <encoding></encoding>
      <mime-type>application/octet-stream</mime-type>
      <data>
       b3JnLmpib3NzLnJlbW90aW5nLkNhbm5vdENvbm5lY3RFeGNlcHRpb246IENhbiBub3QgZ2V0IGNv
       bm5lY3Rpb24gdG8gc2VydmVyLiAgUHJvYmxlbSBlc3RhYmxpc2hpbmcgc29ja2V0IGNvbm5lY3Rp
       [...]
       </data>
    </content>
    <filename>error.txt</filename>
  </schema>
```

27.1.2. Inlining Blobs

There is an option to inline the blob content in the XML file as a Base64 encoded text. This is less optimized but this is the canonic format to export a document data prior to XSD validation of document schemas.

Of course this is less optimized than writing the raw blob data in external files but provides a way to encode the entire document content in a single file and in a well known and validated format.

By default when exporting documents from the repository blobs are not inlined. To activate the inlining option

you must set call the method on the DocumentModelReader you are using to fetch data from the repository:

```
reader.setInlineBlobs(boolean inlineBlobs);
```

27.2. Document Pipe

An export process is a chain of three sub processes:

1. fetching data from repository
2. transforming the data if necessary
3. writing the data to an external system

In the same way an import can be defined as a chain of three sub processes:

1. fetching data from external sources
2. transforming the data if necessary
3. writing the data into the repository

We will name the process chain used to perform imports and exports as a *Document Pipe.*

In both cases (imports and exports) a document pipe is dealing with the same type of objects:

1. A document reader
2. Zero or more document transformers
3. A document writer

So the DocumentPipe will use a reader to fetch data that will be passed through registered transformers and then written down using a document writer.

See the API Examples for examples on how to use a Document Pipe.

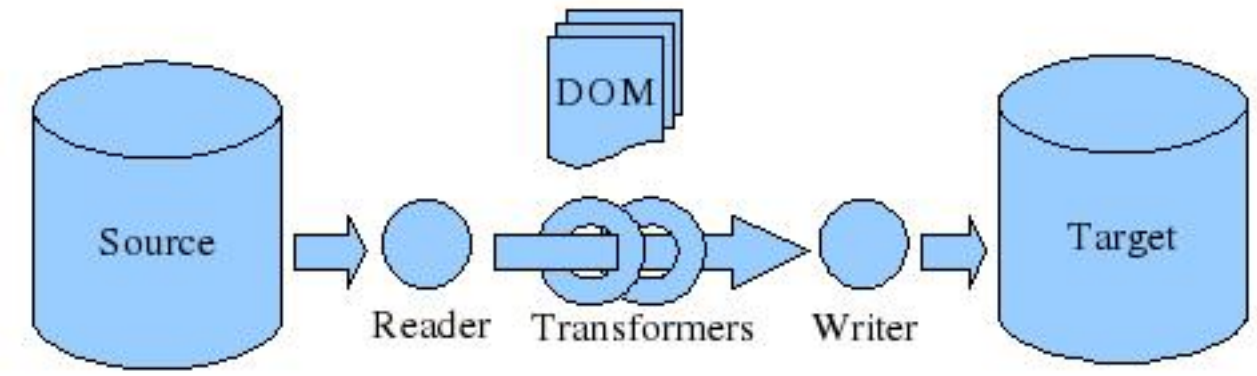

27.3. Document Reader

A document reader is responsible to read some input data and convert it into a DOM representation. The DOM representation is using the format explained in Document XML section. Currently dom4j Documents are used as the DOM objects.

For example a reader may extract documents from the repository and to output it as XML DOM objects. Or it

may be used to read files from a file system and convert them into DOM objects to be able to import them in a Nuxeo repository.

To change the way document are extracted and transformed to a DOM representation you can implement your own Document Reader. Currently Nuxeo provides several flavors of document readers:

1. Repository readers - these category of readers are used to extract data from the repository as DOM objects. All of these readers are extending DocumentModelReader:
   - SingleDocumentReader - this one reads a single document given its ID and export it as a dom4j Document.
   - DocumentChildrenReader - this one reads the children of a given document and export each one as dom4j Document.
   - DocumentTreeReader - this one reads the entire subtree rooted in the given document and export each node in the tree as a dom4j Document.
   - DocumentListReader - this one is taking as input a list of document models and export them as domj Documents. This is useful when wanting to export a search result for example.
2. External readers used to read data as DOM objects from external sources like file systems or databases. The following readers are provided:
   - XMLDirectoryReader - read a directory tree in the format supported by Nuxeo (as described in Export Format section). This can be used to import deflated nuxeo archives or hand created document directories.
   - NuxeoArchiveReader - read Nuxeo EP exported archives to import them in a repository. Note that only zip archives created by nuxeo exporter are supported.
   - ZipReader - read a zip archive and output DOM objects. This reader can read both Nuxeo zip archives and regular zip archives (hand made). Reading a Nuxeo archive is more optimized - because Nuxeo zip archives entries are added to the archive in a predefined order that makes possible reading the entire archive tree on the fly without unziping the content of the archive on the filesystem first. If the zip archive is not recognized as a Nuxeo archive the zip will be deflated in a temporary folder on the file system and the XMLDirectoryReader will be used to read the content.

To create a custom reader you need to implement the interface `org.nuxeo.ecm.core.api.io.DocumentReader`

27.4. Document Writer

A document writer is responsible to write the documents that exit the pipe in a document store. This storage can be a File System, A Nuxeo Repository or any database or data storage as long as you have a writer that supports it.

The following DocumentWriters are provided by Nuxeo:

1. Repository Writers - These ones are writing documents to a Nuxeo repository. They are useful to perform imports into the repository.
   - `DocumentModelWriter` - writes documents inside a Nuxeo Repository. This writer is creating new document models for each one of the imported documents.
   - `DocumentModelUpdater` - writes documents inside a Nuxeo Repository. This writer is updating documents that have the same ID as the imported ones or create new documents otherwise.
2. External Writers - are writers that write documents on an external storage. They are useful to perform exports from the repository.

- `XMLDocumentWriter` - writes a document as a XML file with blobs inlined.
- `XMLDocumentTreeWriter` - writes a list of documents inside a unique XML file with blobs inlined. The document tags will be included in a root tag

```
<documents> .. </documents>
```

- `XMLDirectoryWriter` - writes documents as a folder tree on the file system. To read back the exported tree you may use XMLDirectoryReader
- `NuxeoArchiveWriter` - writes documents inside a Nuxeo azip archive. To read back the archive you may use the NuxeoArchiveReader

To create a custom writer you need to implement the interface `org.nuxeo.ecm.core.api.io.DocumentWriter`

27.5. Document Transformer

Document transformers are useful to transform documents that enter the pipe and before being sent to the writer. This way you can remove, add or modify some properties from the documents, or other information contained by the exported DOM object.

As documents are expressed as XML DOM objects you can also use XSLT transformations inside your transformer.

To create a custom transformer you need to implement the interface `org.nuxeo.ecm.core.api.io.DocumentTransformer`

27.6. API Examples

Performing exports and imports can be done by following these steps:

1. Instantiate a new DocumentPipe:

```
// create a pipe that will process 10 documents on each iteration
DocumentPipe pipe = new DocumentPipeImpl(10);
```

   The page size argument is important when you are running the pipe on a machine different than the one containing the source of the data (the one from where the reader will fetch data). This way you can fetch several documents at once improving performances.

2. Create a new DocumentReader that will be used to fetch data and put it into the pipe. Depending on the data you want to import you can choose between existing DocumentReader implementation or you may write your own if needed:

```
reader = new DocumentTreeReader(docMgr, src, true);
pipe.setReader(reader);
```

   In this example we use a DocumentTreeReader which will read an entire sub-tree form the repository rooted in 'src' document.

   The `docMgr` argument represent a session to the repository, the 'src' is the root of the tree to export and the 'true' flag means to exclude the root from the exported tree.

3. Create a `DocumentWriter` that will be used to write down the outputed by the pipe.

```
writer = new XMLDirectoryWriter(new File("/tmp/export"));
pipe.setWriter(writer);
```

In this example we instantiate a writer that will write exported data onto the file system as a folder tree.

4. Optionally you may add one or more Document Transformers to transform documents that enters the pipe.

```
MyTransformer transformer = new MyTransformer();
pipe.addTransformer(transformer);
```

5. And now run the pipe ...

```
pipe.run();
```

27.6.1. Exporting data from a Nuxeo repository to a Zip archive

```
DocumentReader reader = null;
DocumentWriter writer = null;

try {
  DocumentModel src = getTestWorkspace();
  reader = new DocumentTreeReader(docMgr, root, true);
  writer = new NuxeoArchiveWriter(new File("/tmp/export.zip"));
  // creating a pipe
  DocumentPipe pipe = new DocumentPipeImpl(10);
  pipe.setReader(reader);
  pipe.setWriter(writer);
  pipe.run();
} finally {
  if (reader != null) {
    reader.close();
  }
  if (writer != null) {
    writer.close();
  }
}
```

27.6.2. Importing data from a Zip archive to a Nuxeo repository

```
DocumentReader reader = null;
DocumentWriter writer = null;
try {
  DocumentModel src = getTestWorkspace();
  reader = new ZipReader(new File("/tmp/export.zip"));
  writer = new DocumentModelWriter(docMgr, "import-domain/Workspaces/ws");

  // creating a pipe
  DocumentPipe pipe = new DocumentPipeImpl(10);
  pipe.setReader(reader);
  pipe.setWriter(writer);
  pipe.run();
} finally {
  if (reader != null) {
    reader.close();
  }
  if (writer != null) {
    writer.close();
  }
}
```

27.6.3. Export a single document as an XML with blobs inlined.

```
DocumentReader reader = null;
DocumentWriter writer = null;

try {
  DocumentModel src = getTestWorkspace();
  reader = new SingleDocumentReader(docMgr, src);

  // inline blobs
  ((DocumentTreeReader)reader).setInlineBlobs(true);
  writer = new XMLDocumentWriter(new File("/tmp/export.zip"));

  // creating a pipe
  DocumentPipe pipe = new DocumentPipeImpl();

  // optionally adding a transformer
  pipe.addTransformer(new MyTransformer());
  pipe.setReader(reader);
  pipe.setWriter(writer); pipe.run();

} finally {
  if (reader != null) {
    reader.close();
  }
  if (writer != null) {
    writer.close();
  }
}
```

Chapter 28. Nuxeo Event Service

This chapter presents the event model used in Nuxeo 5.2. For more informations about the differences between this model (5.2-M4+) and the former model used in Nuxeo 5.1, please refer to last section.

28.1. Nuxeo event model

28.1.1. Event

When an event occurs in Nuxeo, an `org.nuxeo.ecm.core.event.Event` object is created and propagated.

`Event` is a lightweight object that consist of :

- a event name
- a timestamp
- some control flags

  flags are used to give informations about the event nature and how it must be processed and forwarded ()
- an `EventContext`

  provides informations on the context assocated to the event when it was fired

28.1.2. EventContext

Each `Event` if always associated to an `EventContext` that holds information about the operation that triggered the event.

Default implementation is very abstract to be very generic and reusable in a lot of situations. Basically en EventContext holds :

- Arguments

  represents arguments of the operation being processed when the event is fired (ordoned list of Serializable Objects)
- Properties

  associates named properties that can be shared between sources of events and listeners

Because the document repository is one of the main sources of events, a `DocumentEventContext` implementation is provided. DocumentEventContext hold :

- The core session used when the event was fired

  (this is the first argument)
- Principal that was logged in when the event was fired

  (this is the second argument)
- The source `DocumenModel`

  (this is the third argument)
- An optional category

(this is a named property)

- An optional comment

  (this is a named property)

EventContext can be used to produce events :

```
DocumentEventContext ctx = new DocumentEventContext(getCoreSession(), getPrincipal(), sourceDocument);
        Event event = ctx.newEvent("MyEvent");
```

28.1.3. EventListener

When an Event is fired, the EventService will call all EventListener in a row.

EventListeners are called synchronously and in an ordered way.

EventListeners have direct access to the original EventContext (CoreSession, User identity ...) and have the possibility to alter this context. Typically an EventListener can intercep all document creation events and automaticall set some fields in the DocumenModel (ex: creation date).

EventListeners can be java classes or scripts.

28.1.4. Transactions and Events

Event firing and EventListeners execution always occur in the same transaction (if any) and in the orginal context.

This means :

- EventListener must be fast

  other wise all transaction may become slow

- EventListener can rollback the current transaction

  either by throwing an unchecked Exception or by setting the `ROLLBACK` flag.

Events that occur in the same transaction are stacked in an `EventBundle` unless they are flagged `INLINE`.

When transaction commits, the associated EventBundle will be fired and stack will be cleaned. In non transactionnal environment, events with the `COMMI` flag will play the role of placeholders.

Typically, if no JTA environment is available, all call to CoreSession.save() will fire the event "SAVE" that will do a pseudo-commit on stacked events.

28.1.5. EventBundle

EventBundle represent the stack of events that have occured in the same transaction.

And end of transaction, the bundle will be fired.

28.1.6. PostCommitEventListener

PostCommitListeners are notified when `EventBundle` are fired.

There are 2 types of PostCommitEventListeners :

- Synchronous PostCommitEventListener

  Execution occurs after the transaction commits, but before the call has returned to the client.

- Asynchronous PostCommitEventListener

  Execution occurs after the transaction commits, and after the call has returned to the client.

It is important to understand that in both cases, each PostCommitEventListener is executed in a separated transaction. Each PostCommitEventListener may commit or rollback without affecting other listeners and the main transaction. In the case of asynchronous PostCommitEventListener, the EventContext is not exactly the same as the orginal one, this is a recustructed EventContext :

- Security context is switched to a System login

  (but informations about the orginal Principal is conserved)

- The CoreSession that may be associated to EventContext is a System one

  (since orginal user's CoreSession may have bee closed)

- Associated DocumentModels are "re-fetched" from the CoreSession

28.2. Using Events

28.2.1. Firing Events

Firing an event is very simple :

```
// get EventProdicer Service
EventProducer evtProducer = Framework.getService(EventProducer.class);

// prepare EventContext properties
Map<String, Serializable> props = new HashMap<String, Serializable>();
props.put("myinfo", "foo-bar");

// create simple EventContext
EventContext ctx = new InlineEventContext(principal, props);

// create the event from the context
Event event = ctx.newEvent(eventId);

// fire the event
evtProducer.fireEvent(event);
```

This exemple demontrate firing of a simple event that won't be stacked in a bundle (not tied to any transaction).

```
// Create EventContext bound to a Document related operation
DocumentEventContext ctx = new DocumentEventContext(session,principal, myDoc);
ctx.setCategory("MyEventCategory");
ctx.setComment("MyComment");

// prepare EventContext properties
Map<String, Serializable> props = new HashMap<String, Serializable>();
props.put("myinfo", "foo-bar");
ctx.setProperties(props);

// fire the event
producer.fireEvent(event);
```

This exemple demontrate firing of a Document related event.

28.2.2. Contributing an EventListener

Event listener are contributed via `listener` extension point of

`org.nuxeo.ecm.core.event.EventServiceComponent`.

Contributed listener can be :

- A java class

  must implement the `EventListener` interface

- A script

Here is an example of a contributed EventListener is Java :

```
<extension target="org.nuxeo.ecm.core.event.EventServiceComponent" point="listener">
    <listener name="myListener"
              async="false"
              postCommit="false"
              class="com.myproject.listener.MySyncEventListener"
              priority="140">
    </listener>
</extension>
```

The `async` and `postCommit` attribute are not necessary since the java class interface already provides this information.

Here is an example of a contributed EventListener is Java :

```
<extension target="org.nuxeo.ecm.core.event.EventServiceComponent" point="listener">
    <listener name="myScriptListener"
              async="false"
              postCommit="false"
              script="script/listener.groovy"
              priority="145">
    </listener>
</extension>
```

In this case, the `async` and `postCommit` attribute are necessary. The groovy script will receive as input the `Event`.

28.2.3. Contributing a PostCommitEventListener

PostCommitEventListener are contributed in the same way that EventListener, using the same extension point. The EventService will determine that the listener is a PostCommitEventListener based on :

- The interface of the contributed java class

  Interface `PostCommitEventListener`

- the async and postcommit attribute.

Exactly like for synchronous EventListener, PostCommitEventListener can be a Java class or a script.

On contrary of synchronous EventListeners, PostCommitEventListener will receive an EventBundle (containing all events associated to a transaction), instead of a single Event.

Here is an example of a contributed PostCommitEventListener is Java :

```
<extension target="org.nuxeo.ecm.core.event.EventServiceComponent" point="listener">
    <listener name="myAsyncListener"
              async="false"
              postCommit="true"
              class="com.myproject.listener.MyAsyncEventListener"
              priority="140">
    </listener>
</extension>
```

The `postCommit` flag is not really useful since the java interface already defines the listener as PostCommit. The `async` flag is used to define if the PostCommitListener should be processed before or after returning the call to the client.

28.3. JMS and Nuxeo Events

28.3.1. JMS integration

JMS can be used to relay events from one JVM to other JVMs hosting Nuxeo Components. If a multi-server deployment is used, it allow to have an async PostCommitEventListener running on one JVM even if the event was fired on another JVM.

In multi-JVM deployment scenarios, the EventService must be deployed on each JVM.

Separated instances of the EventService will be linked via the JMS bus.

The bridge is done with 2 additionnal bundles :

- `nuxeo-core-event-jms`

  PostCommitEventListener that forwards EventBundle to a JMS Topic.

- `nuxeo-platform-event-dispatcher`

  MessageDrivenBean that intercepts JMS message, recustructuc the EventBundle and associated context and fire the EventBundle in the local EventService.

In order to avoid any loop or duplicated EventListener execution :

- `nuxeo-core-event-jms` never process EventBundle that have already been relayed via JMS
- `nuxeo-platform-event-dispatcher` only process EventBundle comming from another JMV

28.3.2. Enabling JMS bridge

In default mono-JVM Nuxeo Bundle, the JMS bridge is not deployed and not activated.

In order to activate the JMS bridge you have to deploy `nuxeo-core-event-jms` and `nuxeo-platform-event-dispatcher`.

If you want to force JMS usage in mono-JVM deployment, you can add org.nuxeo.ecm.event.forceJMS=true in nuxeo.properties.

28.3.3. From 5.1 event model to 5.2

28.3.3.1. Main differences

The API for sending an event is now far more simple in the new model:

- Only one API in 5.2
- There was 2 APIs in the 5.1

  One for Core Events and one for JMS events

Writing an asynchronous eventListener is also simpler :

- no need a write a MessageDrivenBean
- no need to handle Authentication and Repository initialization
- no need to handle transactions

One of the key benefit of the refactoring, is that JMS is no longer a required dependency : you can have asynchronous listeners without JMS. This means :

- we can embed more services in a WebEngine/Jetty package
- you can now easily test your eventListeners in JUnit without needing JMS/MDB setup

28.3.3.2. About compatibility

A bundle `nuxeo-core-event-compat` is provided. When deployed, this bundle will :

- to run "old style" CoreEventListeners
- to enable JMS forwarding on the 5.1 JMS Topic

This means that with this compatibility module, you can run old CoreEventListeners and MessageDrivenBeans.

For sending JMS events using the old API, you will need to deploy `nuxeo-platform-events-api`

Chapter 29. Experimental Topics

29.1. Introduction

This chapter is provided to document experimental features that are being introduced in the current development branch (Nuxeo 5.2-SNAPSHOT).

This content is likely to move to other places after some time.

Note

Your feedback is welcome! If you have use cases and want to help with the development of these technologies, please join the ECM mailing list or the forum.

29.2. Runtime Support for Scripting Languages

29.2.1. Introduction

Preliminary, yet powerful, support for scripting in Nuxeo Runtime has been recently added. This makes scripting available from all Nuxeo's platform. Thanks to this new feature, you can easily uses script from your custom components. This can be very useful for a lot of use cases, like:

- dynamic rules (scripting language as DSLs)
- easily modifiable behaviors
- light and powerful configuration / customization
- etc.

Scripts have access to the whole API thanks to Java scripting integration (JSR-223).

Moreover, scripts can also be run remotely thanks to the Nuxeo Runtime command line. This allows you to create a script on your administration machine, launch it on the remote platform and get the result back. It makes scripting a killer-feature for administration scripts (ex: expire content, bulk content modification, bulk refactoring of the content repository layout, etc.).

29.2.2. Supported languages

I've just integrated scripting support through JSR 223 in nuxeo. This was integrated as a new project `nuxeo-runtime-scripting` which is in SVN but it is was not yet added in the `nuxeo.ear` build (neither in runtime SVN module).

For now only these scripting engines have been integrated:

- Jexl
- JRuby
- Jython
- Groovy
- JavaScript (Rhino)

If needed more will be added later (like PHP for example).

29.2.3. Running a Script

You can run scripts in nuxeo in 3 different ways:

- Put the script inside the `nuxeo.ear/script` directory (you should define this directory through a runtime variable). Then from the Java code you can do:

```
Framework.getService(ScriptingService.class).eval("my_script.js");
```

  where `my_script.js` is the script path relative to the script directory.

  Or you can use the JSR 322 API:

```
ScriptEngineManager factory = new ScriptEngineManager();
ScriptEngine engine = factory.getEngineByName("js");
engine.eval("absolute/path/to/my_script.js");
```

- Let the script inside your `.jar` and register it under using name as a script component. Then you can run the script as follow:

```
Framework.getService(ScriptingService.class).getScript("myScript").eval();
```

  This method chaches the compiled script so it is only supported for languages that support compilation (currently all the engines that comes in Nuxeo).

- Run a script from remote :)

  This can be used for debug, testing or administration. You write a script locally then you run it against a remote Nuxeo EP server. The script will be send to the server and executed on the server then the server will return the result (including STDOUT and STDERR) to the client.

  For security reason this feature can be disabled using a runtime property on the server.

  Here is an example on how you can run a remote script:

```
ScriptingClient client = new ScriptingClient("localhost", 62474);
URL src = RemoteTest.class.getClassLoader().getResource("test.js");
RemoteScript script = client.loadScript(src);
script.eval();
```

  For the following JavaScript script:

```
importPackage(java.lang);
importPackage(org.nuxeo.runtime);
importPackage(org.nuxeo.runtime.api);
importPackage(org.nuxeo.runtime.model);

runtime = Framework.getRuntime();
name = runtime.getName();
version = runtime.getVersion();
desc = runtime.getDescription();
println("Remote runtime: "+name+" v."+version);
println(desc);
println("----------------------------------------");
println("Registered components:");
println("----------------------------------------");
regs = runtime.getComponentManager().getRegistrations();
for (var i=0, size=regs.size(); i<size; i++){
    println(regs.get(i).getName());
}
```

  The following will be printed on the STDOUT of the client:

```
Remote runtime: OSGi NXRuntime v.1.4.0
OSGi NXRuntime version 1.4.0
----------------------------------------
Registered components:
----------------------------------------
```

```
service:org.nuxeo.ecm.core.api.DocumentAdapterService
service:org.nuxeo.ecm.core.repository.RepositoryService
service:org.nuxeo.runtime.remoting.RemotingService
service:org.nuxeo.runtime.EventService
service:org.nuxeo.ecm.platform.login.LoginConfig
...
```

So, this new feature can be used to write pure script based Nuxeo components. Also in future I will try to configure Tomcat to be able to run scripts inside servlets. This means to be able to write web pages in PHP or other supported language for Nuxeo EP ;-)

If you're willing to leverage these new scripting features in your own applications and explore new ways to use scripting to make balance "enterprise" with "agile" in the ECM field, please join the list or the forum and start contributing.

Part IV. SOA, Web Services and various integration solutions

Chapter 30. The Nuxeo Restlet API

Nuxeo integrates the Restlet framework to provide an easy way to contribute new REST API on top of the platform.

30.1. Restlet Integration

The REST API provides an easy way to call Nuxeo services from an external application. Even if REST is a very simple concept, we choose to leverage an existing REST framework to provide an REST API on top of Nuxeo. The selected framework is Restlets (http://www.restlet.org/) that provides a lightweight and flexible REST framework. The Nuxeo REST integration API (`org.nuxeo.ecm.platform.ui.web.restAPI`) provides:

- A runtime service to contribute new restlets
- Base classes for new restlets
- A main servlet that handles the routing between restlets
- An integration with the Nuxeo API and Seam context

To implement a restlet you simply have to implement the `Restlet` interface.

30.1.1. Restlet types in Nuxeo 5

Nuxeo 5 defines 3 different types of restlets :

- Stateless restlets

  No integration with Seam and no state management. This is the original logic of Restlet.

  You can use `BaseStalessNuxeoRestlet` as base class that provides helpers for accessing main services (like the repository).

- Seam-aware restlets

  For restlet declared as Seam-aware, the Nuxeo Restlet servlet initializes the Seam context before executing the restlet. Thanks to this initialization your restlet can use injection (`@In`) to access the Seam context. This solution gives you the possibility of using existing Seam components. You don't have to use Service Platform API to access the service since you can access Seam delegates for that.

  You can use `BaseNuxeoRestlet` that provides helper API for error handling, security and URL management.

- Conversation-aware restlets

  The are Seam restlets that are tied to a Seam conversation. For these restlets, the Seam context initialization is done in order to setup the current Conversation. Conversational restlets must be called with a `conversationId` as parameter.

  Conversational restlet can be used if you need to access the current Seam context associated with a browser session. Typically this is what is used by the Firefox helper to upload files.

  You can use `BaseNuxeoRestlet` that provides helper API for error handling, security and URL management.

30.1.2. Restlet URL and parameters mapping

All restlets are bound to one or more URL patterns. These URLs patterns are used by Nuxeo `RestletServlet` to determine which restlet needs to receive the call.

When defining URLs you can use {} to have parts of the URL that will be converted as parameters.

For example :

```
/{repo}/{docid}/{filename}/upload
```

will define a URL pattern with 4 parts and you will have access from withing your code to 3 parameters: repo, docId and filename.

These parameters can be used via Restlet API :

```
req.getAttributes().get("repo");
```

You can also access the standard GET/POST parameters via

```
req.getResourceRef().getQueryAsForm().getFirstValue("SomeParameter");
```

30.1.3. Contributing a new restlet

Contributing a new restlet is quite simple.

The first thing to do is to write a new restlet: you can either implement the `Restlet` Interface "by hand" or just inherit from `BaseNuxeoRestlet` or `BaseStatelessNuxeoRestet`.

Once your class is written, you need to contribute to the restlets extension point exposed by `org.nuxeo.ecm.platform.ui.web.restAPI.service.PluggableRestletService`.

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.platform.ui.web.restAPI.contrib">

  <extension
      target="org.nuxeo.ecm.platform.ui.web.restAPI.service.PluggableRestletService"
      point="restlets">
    <restletPlugin
        name="upload"
        class="org.nuxeo.ecm.platform.ui.web.restAPI.UploadRestlet"
        enabled="true"
        useSeam="true"
        useConversation="false">
      <urlPatterns>
        <urlPattern>/{repo}/{docid}/{filename}/upload</urlPattern>
      </urlPatterns>
   </restletPlugin>
   ...
```

The `useSeam` and `useConversation` flags define how the Nuxeo Restlet servlet will handle the call.

30.2. Nuxeo default restlets

Nuxeo comes by default with very simple restlet that can be seen as samples.

30.2.1. Browse restlet

The Browse restlet is a simple way to navigate into the repository via REST.

A typical call to list content of default repository would be :
http://127.0.0.1:8080/nuxeo/restAPI/default/*/browse

If you need to browse the document 95ce52b2-6959-4afa-bc63-396096b376b4 a typical call would be :
http://127.0.0.1:8080/nuxeo/restAPI/default/95ce52b2-6959-4afa-bc63-396096b376b4/browse

Typical output for this restlet would be:

```
<document title="2fbf878d-9c2f-42c6-acbb-ea339ce15615"
          type="Root"
```

```
          id="2fbf878d-9c2f-42c6-acbb-ea339ce15615"
          url="/default/2fbf878d-9c2f-42c6-acbb-ea339ce15615">
  <document title="Default domain"
            type="Domain"
            id="95ce52b2-6959-4afa-bc63-396096b376b4"
            url="/default/95ce52b2-6959-4afa-bc63-396096b376b4"/>
</document>
```

This restlet uses Seam in order to have documentManager injected (this is not a need but rather a simple way of accessing the repository without using the Service API)

Browse restlet registration looks like this:

```
<restletPlugin
    name="browse"
    class="org.nuxeo.ecm.platform.ui.web.restAPI.BrowseRestlet"
    enabled="true"
    useSeam="true">
  <urlPatterns>
    <urlPattern>/{repo}/{docid}/browse</urlPattern>
  </urlPatterns>
</restletPlugin>
```

30.2.2. Export restlet

This restlet is a simple REST frontend on top of IO core service. This can be used to export a document or a document tree as XML.

A typical call would be :

- http://127.0.0.1:8080/nuxeo/restAPI/default/2fbf878d-9c2f-42c6-acbb-ea339ce15615/export

  to export the document 2fbf878d-9c2f-42c6-acbb-ea339ce15615 as XML

- http://127.0.0.1:8080/nuxeo/restAPI/default/2fbf878d-9c2f-42c6-acbb-ea339ce15615/export?format=ZIP

  to export the document 2fbf878d-9c2f-42c6-acbb-ea339ce15615 as a ZIP archive

- http://127.0.0.1:8080/nuxeo/restAPI/default/2fbf878d-9c2f-42c6-acbb-ea339ce15615/exportTree

  to export the document 2fbf878d-9c2f-42c6-acbb-ea339ce15615 and all its children as a Zip Archive

- http://127.0.0.1:8080/nuxeo/restAPI/default/2fbf878d-9c2f-42c6-acbb-ea339ce15615/exportTree?format=XML

  to export the document 2fbf878d-9c2f-42c6-acbb-ea339ce15615 and all its children as XML

This restlet uses Seam in order to have documentManager injected (this is not a need but rather a simple way of accessing the repository without using the Service API)

Export restlet registration looks like this:

```
<restletPlugin
    name="export"
    class="org.nuxeo.ecm.platform.ui.web.restAPI.ExportRestlet"
    enabled="true"
    useSeam="true">
  <urlPatterns>
    <urlPattern>/{repo}/{docid}/export</urlPattern>
    <urlPattern>/{repo}/{docid}/exportSingle</urlPattern>
    <urlPattern>/{repo}/{docid}/exportTree</urlPattern>
  </urlPatterns>
</restletPlugin>
```

30.2.3. Lock restlet

This restlet provide a REST API for Lock Management.

A typical call would be:

- GET http://127.0.0.1:8080/nuxeo/restAPI/default/2fbf878d-9c2f-42c6-acbb-ea339ce15615/Locking/status

  to get Lock status of the document 2fbf878d-9c2f-42c6-acbb-ea339ce15615

- GET http://127.0.0.1:8080/nuxeo/restAPI/default/2fbf878d-9c2f-42c6-acbb-ea339ce15615/Locking/lock

  or

  LOCK http://127.0.0.1:8080/nuxeo/restAPI/default/2fbf878d-9c2f-42c6-acbb-ea339ce15615/Locking

  to Lock status the document 2fbf878d-9c2f-42c6-acbb-ea339ce15615

- GET http://127.0.0.1:8080/nuxeo/restAPI/default/2fbf878d-9c2f-42c6-acbb-ea339ce15615/Locking/unlock

  or

  UNLOCK http://127.0.0.1:8080/nuxeo/restAPI/default/2fbf878d-9c2f-42c6-acbb-ea339ce15615/Locking

  to Unlock status the document 2fbf878d-9c2f-42c6-acbb-ea339ce15615

This restlet is stateless and uses Nuxeo Service API

Lock restlet registration looks like this

```
<restletPlugin
    name="locking"
    class="org.nuxeo.ecm.platform.ui.web.restAPI.LockingRestlet"
    enabled="true"
    useSeam="false"
    useConversation="false">
  <urlPatterns>
    <urlPattern>/{repo}/{docid}/Locking</urlPattern>
  </urlPatterns>
</restletPlugin>
```

30.2.4. Plugin upload restlet

This restlet provides a REST API for file upload that is used by the Firefox helper.

A typical call would be:

- POST http://127.0.0.1:8080/nuxeo/restAPI/default/e5125b5e-8b9e-43bd-8959-7e7e5caf2a1b/pluginUpload/myfolder/myfile

  to upload a file

This restlet uses Seam conversation.

pluginUpload restlet registration looks like this

```
<restletPlugin
    name="pluginUpload"
    class="org.nuxeo.ecm.platform.ui.web.restAPI.PluginUploadRestlet"
    enabled="true"
    useSeam="true"
    useConversation="true">
  <urlPatterns>
    <urlPattern>/{repo}/{docid}/pluginUpload</urlPattern>
  </urlPatterns>
</restletPlugin>
```

30.3. Nuxeo RestPack

The RestPack is an additional component for Nuxeo EP that provides additional restlets on top of Nuxeo Services.

The primary goal of the RestPack is to provide the needed RestAPI to build simple JSR168 portlets (like search portlets) that communicate with Nuxeo via HTTP/XML.

30.3.1. Installing the RestPack

You can download the RestPack from Nuxeo download site or build it from source.

To deploy the module, just copy the jar file in nuxeo.ear/system folder and restart you JBoss

30.3.2. Restlets included in the RestPack

30.3.2.1. Vocabulary Restlet

The Vocabulary restlet exports as XML the content of a given vocabulary

Sample call :

GET http://127.0.0.1:8080/nuxeo/restAPI/vocabulary/{vocName}/

parameters

- {vocName}

  Name of the vocabulary to export, must be last part of QueryPath

- lang

  GET (QueryString) parameter used to set language used to generate the labels

Sample calls

- http://127.0.0.1:8080/nuxeo/restAPI/vocabulary/subject?lang=en

```
<entries>
<entry id="arts" label="label.directories.subject.arts" translatedLabel="Arts"/>
<entry id="business" label="label.directories.subject.business" translatedLabel="Business"/>
<entry id="computers" label="label.directories.subject.computers" translatedLabel="Computers"/>
...
```

- http://127.0.0.1:8080/nuxeo/restAPI/vocabulary/continent_country?lang=en

```
<entries>
<entry id="africa" label="label.directories.continent.africa" translatedLabel="Africa">
<entry id="Algeria" label="label.directories.country.Algeria" translatedLabel="Algeria" parent="africa"/>
<entry id="Angola" label="label.directories.country.Angola" translatedLabel="Angola" parent="africa"/>
<entry id="Benin" label="label.directories.country.Benin" translatedLabel="Benin" parent="africa"/>
<entry id="Botswana" label="label.directories.country.Botswana" translatedLabel="Botswana" parent="africa"/>
<entry id="Burkina_Faso" label="label.directories.country.Burkina_Faso" translatedLabel="Burkina Faso" parent="a
<entry id="Burundi" label="label.directories.country.Burundi" translatedLabel="Burundi" parent="africa"/>
<entry id="Cameroon" label="label.directories.country.Cameroon" translatedLabel="Cameroon" parent="africa"/>
...
```

30.3.2.2. Syndication Restlet

Export as RSS or ATOM the list of children document of the targeted container

Sample call:

http://127.0.0.1:8080/nuxeo/restAPI/{repoId}/{docId}/{format}

Parameters:

- {repoId}

  Name of the target repository (use default on a standard installation)

- {docId}

  DocumentRef of the target container (use the docId present in the standard nuxeo URL when browsing the webapp)

- {format}

  Defines the syndication format: rss or atom

30.3.2.3. Workflow Tasks Restlet

Export user's workflow tasks as XML or ATOM

Sample call:

http://127.0.0.1:8080/nuxeo/restAPI/workflowTasks/{repoId}/?format=XML/ATOM

Parameters:

- {repoId}

  Name of the target repository (use default on a standard installation)

- format

  Defines the export format: xml or atom

Sample call:

http://127.0.0.1:8080/nuxeo/restAPI/workflowTasks/default/?format=xml

```
<nxt:tasks>
 <nxt:category category="None">
  <nxt:task name="review" workflowType="document_review_approbation" author="Administrator" startDate="2007-10-0
 </nxt:category>
</nxt:tasks>
```

30.3.2.4. Query for Workflow Tasks Restlet

Export user's workflow tasks that match the query as XML or ATOM

Sample call:

http://127.0.0.1:8080/nuxeo/restAPI/queryForWorkflowTask/{repoId}/?format=XML/ATOM&workflowRequest=author&

Parameters:

- {repoId}

  Name of the target repository (use default on a standard installation)

- format

Defines the export format: xml or atom

- workflowRequest

  Defines the workflow item attribute to compare for the query. Please refer to the WorkItem Filter.

- comparisonType

  Defines the type of comparison. Please refer to the WorkItem Filter.

- workItemFromUser

  If is true add filter that match WorkItem initiated by the user

- canManage

  If true add filter that match workItems that the participant has a direct action to perform

Sample call:

http://127.0.0.1:8080/nuxeo/restAPI/workflowTasks/default/?format=xml

```
<nxt:tasks>
 <nxt:category category="None">
  <nxt:task name="review" workflowType="document_review_approbation" author="Administrator" startDate="2007-10-0
 </nxt:category>
</nxt:tasks>
```

http://localhost:8080/nuxeo/restAPI/queryForWorkflowTask/default/?workflowRequest=author&comparisonType=0&work

Get work item tasks that the user initiate and is pending for the next user.

30.3.2.5. QueryModel Restlet

Execute a search via QueryModel and returns DocumenList as XML/RSS/ATOM/JSON

Sample call :

http://127.0.0.1:8080/nuxeo/restAPI/execQueryModel/{queryModelName}

Parameters:

- {queryModelName}

  Name of the QueryModel to execute

- format

  Defines the export format: XML, Atom, RSS or JSON. This parameter is set as a QueryString parameter

- page

  Defines the page number you want in the result. This parameter is set as a QueryString parameter

- ascending

  Defines ordering ascending: true/false. This parameter is set as a QueryString parameter

- criteria

  Defines ordering columns used for sorting. This parameter is set as a QueryString parameter

- columns

  Defines the columns you want to be included in the resultset. This parameter is set as a QueryString parameter

  This parameter is a simple string containing schema.fieldName tokens separated by "," . The only special field is url. Default value is : dublincore.title,dublincore.description,url

- QueryModel parameters

  For stateless QueryModels, you have to specify the parameters via QP1, QP2 ... (only ordering is important)

  The value $USER is automatically replaced by the name of the current User :

  http://127.0.0.1:8080/nuxeo/restAPI/execQueryModel/USER_DOCUMENTS?QP1=$USER

  For stateful QueryModels, you have to specify the parameters as fieldName=FieldValue

Sample call:

http://127.0.0.1:8080/nuxeo/restAPI/execQueryModel/USER_DOCUMENTS?QP1=$USER&format=JSON

```
[
 {
  "title": "nouveau-fichier",
  "description": null,
  "url": "nxdoc/default/452c122d-07de-422b-b448-b0fef9534a62/view_documents",
  "id": "452c122d-07de-422b-b448-b0fef9534a62"
 },
 {
  "title": "Setup",
  "description": null,
  "url": "nxdoc/default/ae03f7bf-9967-4b5b-b37b-807fd40a6ec7/view_documents",
  "id": "ae03f7bf-9967-4b5b-b37b-807fd40a6ec7"
 },
 {
  "title": "cps",
  "description": null,
  "url": "nxdoc/default/2bad93ca-188f-4ea0-a585-9540a6ed6581/view_documents",
  "id": "2bad93ca-188f-4ea0-a585-9540a6ed6581"
 },
 {
  "title": "testMe",
  "description": null,
  "url": "nxdoc/default/e4de81a9-95e8-49ff-9146-e020f99b8bb8/view_documents",
  "id": "e4de81a9-95e8-49ff-9146-e020f99b8bb8"
 } ... ]
```

http://127.0.0.1:8080/nuxeo/restAPI/execQueryModel/USER_DOCUMENTS?QP1=$USER&format=XML&page=1

```
<results>
<pages pages="3" pageNumber="1"/>
<document id="a3154f03-6baa-4d7d-8bac-579d52a8d304" title="ssl-uil2-service"
    url="nxdoc/default/a3154f03-6baa-4d7d-8bac-579d52a8d304/view_documents"/>
<document id="e2b26d9a-9140-44f1-ae23-4ad7866911c0" title="build"
    url="nxdoc/default/e2b26d9a-9140-44f1-ae23-4ad7866911c0/view_documents"/>
<document id="bf7de6df-bca1-4d29-8e26-5019ced696fd" title="jboss-service"
    url="nxdoc/default/bf7de6df-bca1-4d29-8e26-5019ced696fd/view_documents"/>
<document id="439a11e8-642d-41f7-ae0f-4ef4d7303d79" title="postgres-jdbc2-service"
    url="nxdoc/default/439a11e8-642d-41f7-ae0f-4ef4d7303d79/view_documents"/>
<document id="467745db-565d-4350-953f-b2fa03733406" title="oil-service"
    url="nxdoc/default/467745db-565d-4350-953f-b2fa03733406/view_documents"/>
<document id="44a2441e-e265-4fee-ad17-9ec24245cccf" title="jbossmq-state"
    url="nxdoc/default/44a2441e-e265-4fee-ad17-9ec24245cccf/view_documents"/>
<document id="8d7e3077-a3be-4d6c-806c-aa97816821e2" title="oracle-jdbc2-service"
    url="nxdoc/default/8d7e3077-a3be-4d6c-806c-aa97816821e2/view_documents"/>
<document id="b431ad04-3439-441e-826b-48b5f0f0c1c8" title="as400-jdbc2-service"
    url="nxdoc/default/b431ad04-3439-441e-826b-48b5f0f0c1c8/view_documents"/>
<document id="e146c6b2-a315-4788-9bbe-5f1fc79c18ac" title="mssql-jdbc2-service"
    url="nxdoc/default/e146c6b2-a315-4788-9bbe-5f1fc79c18ac/view_documents"/>
<document id="b109a672-5ebd-4a01-8a86-79e2430c2f07" title="file-state-service"
    url="nxdoc/default/b109a672-5ebd-4a01-8a86-79e2430c2f07/view_documents"/>
</results>
```

http://127.0.0.1:8080/nuxeo/restAPI/execQueryModel/USER_DOCUMENTS?QP1=$USER&format=XML&page=1&colu

```
<results>
<pages pages="3" pageNumber="1"/>
<document id="a3154f03-6baa-4d7d-8bac-579d52a8d304" icon="/icons/note.gif" title="ssl-uil2-service"/>
<document id="e2b26d9a-9140-44f1-ae23-4ad7866911c0" icon="/icons/note.gif" title="build"/>
<document id="bf7de6df-bca1-4d29-8e26-5019ced696fd" icon="/icons/note.gif" title="jboss-service"/>
<document id="439a11e8-642d-41f7-ae0f-4ef4d7303d79" icon="/icons/note.gif" title="postgres-jdbc2-service"/>
<document id="467745db-565d-4350-953f-b2fa03733406" icon="/icons/note.gif" title="oil-service"/>
<document id="44a2441e-e265-4fee-ad17-9ec24245cccf" icon="/icons/note.gif" title="jbossmq-state"/>
<document id="8d7e3077-a3be-4d6c-806c-aa97816821e2" icon="/icons/note.gif" title="oracle-jdbc2-service"/>
<document id="b431ad04-3439-441e-826b-48b5f0f0c1c8" icon="/icons/note.gif" title="as400-jdbc2-service"/>
<document id="e146c6b2-a315-4788-9bbe-5f1fc79c18ac" icon="/icons/note.gif" title="mssql-jdbc2-service"/>
<document id="b109a672-5ebd-4a01-8a86-79e2430c2f07" icon="/icons/note.gif" title="file-state-service"/>
</results>
```

30.4. Nuxeo WebEngine Restlets

WebEngine is a new way to develop Restlets inside Nuxeo using the new specification JAX-RS.

Several restlets powered by WebEngine are already available, lets look at http://doc.nuxeo.org/xwiki/bin/view/Main/WebEngineTutorial to learn how to use them.

Chapter 31. Nuxeo HTTP client

Nuxeo HTTP client is a simple helper library to encapsulate restlets calls to Nuxeo platform

Even if you can easily use restlet by directly manipulating HTTP Request (this is one of the purpous of REST), if you use Java on the client side, you may choose to use Nuxeo HTTP client.

The client lib provides two main features:

- Encapsulate restlet HTTP client library
- Encapsulate Nuxeo authentication

31.1. HTTP Client Library

The library mainly encapsulate the Restlet Client API.

The main service object is a `NuxeoServer` that provides attribute configuration and call methods:

```
NuxeoServer nxServer = new NuxeoServer("http://127.0.0.1:8080/nuxeo");

nxServer.setAuthType(NuxeoServer.AUTH_TYPE_BASIC);
nxServer.setBasicAuthentication("Administrator", "Administrator");

List<String> pathParams = Arrays.asList("vocabulary", "country");

Representation res = nxServer.doRestletGetCall(pathParams, null);
```

31.2. HTTP client authentication

For authentication, the HTTO client library proposes 2 implementations:

- Basic Authentication

  Classic Web Authentication

- Shared Secret Authentication

  Designed to be able to impersonate calls

The shared Secret Authentication depends on an additionnal authentication plugin that needs to be deployed on Nuxeo side: `nuxeo-platform-login-portal-sso`.

This authentication system is based on a shared secret between the client and Nuxeo server: you need to configure this shared secret in the configuration file of the server side module, and also to pass this secret to the client http lib. Thanks to this shared secret the client will send the login name and a digest token that will be used to execute the request on the behalf of the login user

A typical use case is a JSR 168 portlet that fetches data from Nuxeo EP. The data retrieval must be done on behalf of the connected user (`request.getPrincipal()`). This allows a portlet to display user's workspaces list, or last documents without the portlet having to know the password of the user.

The authentication token sent between the client is based on the shared secret, the user login, a random data and a timestamp. Althought this should be secure enought for most needs, this trusted communication between a client application and a Nuxeo server should not be done on a HTTP connection that uses public Internet.

Here is a sample call:

```
NuxeoServer nxServer = new NuxeoServer("http://127.0.0.1:8080/nuxeo");

nxServer.setAuthType(NuxeoServer.AUTH_TYPE_SECRET);
```

```
nxServer.setSharedSecretAuthentication("JDoe","nuxeo5secretkey");

List<String> pathParams = Arrays.asList("execQueryModel","USER_DOCUMENTS");

Map<String, String> queryParams = new HashMap<String, String>();

queryParams.put("QP1", "$USER");
queryParams.put("format", "JSON");
Representation res = nxServer.doRestletGetCall(pathParams, queryParams);
```

Chapter 32. Web services

You will find the list of deployed J2EE compatible web services on your host at the following location: http://localhost:8080/jbossws. From there you can view all the web services descriptions (`WSDL` files).

32.1. Audit web service

This web service enable to query all logged events.

TODO: verbose

32.2. Remoting web service

TODO: verbose

32.3. Indexing gateway service

Nuxeo EP provides a WebService API that enables an external indexer to index the documents that are inside the Nuxeo repository. In this scenario, Nuxeo keeps it's internal indexing backend, and may contiue to use it for search.

This service is provided as an addon called `nuxeo-platform-indexing-gateway` that need to be deployed in the `nuxeo.ear/plugins` directory.

The data exchange between the external indexer and Nuxeo is done in a pull mode: the external indexer queries Nuxeo for new documents to index.

Here is a simple summary of the responsibilities between the external indexer and Nuxeo server:

Externalindexer:

- query nuxeo to know what document may be indexed
- query nuxeo to get document informations
- manages indexing configuration

Nuxeo:

- provides an API to browse repository
- provides an API to know what documents needs to be indexed
- may include some transformation of data before exportation to make
- external indexer easier (like ACP/ACL/ACE preprocessing) may notify external indexer that some data need to be (re)indexed

32.4. Metro web services

The WS stack used is Metro 1.4. The stack is deployed by default by running "ant patch" in the nuxeo sources directory. It also removes the default JBoss WS stack. Some examples on how to set up such web services can be found at http://hg.nuxeo.org/sandbox/nuxeo-platform-webservice-example/

The README.txt should be a good start.

Chapter 33. Nuxeo JSR 168 Integration

33.1. Overview

33.2. Testing Nuxeo Portlets

33.2.1. Prerequisites

We will assume that a JBoss is installed in `/opt/jboss`, and a Nuxeo platform is deployed on it.

33.2.1.1. Install the restPack project

The nuxeo-platform-restPack project contains some useful restlets, including the *sample* restlet used in the following project. You need to install it in your deployed Nuxeo EP.

Checkout the sources of the `nuxeo-platform-restPack` project:

```
svn co https://svn.nuxeo.org/nuxeo/nuxeo-addons/nuxeo-platform-restPack/trunk nuxeo-platform-restPack
```

Then, go into the newly created folder, build the project and copy the new .jar:

```
mvn clean package
cp target/nuxeo-platform-restPack-VERSION.jar /opt/jboss/server/default/deploy/nuxeo.ear/system
```

Restart JBoss

33.2.1.2. Authentication

To enable the authentication between the portlet and the Nuxeo server, you need to install the `nuxeo-platform-login-portal-sso` project in your deployed Nuxeo EP.

Checkout the sources of the `nuxeo-platform-login-portal-sso` project:

```
svn co https://svn.nuxeo.org/nuxeo/nuxeo-addons/nuxeo-platform-login-portal-sso/trunk \
nuxeo-platform-login-portal-sso
```

To configure the authentication, edit the `Sample-Portal-SSO-descriptor-bundle.xml` file located in the `Sample` folder. Then, copy it in the `config` folder of your deployed Nuxeo EP.

```
cp Sample/Sample-Portal-SSO-descriptor-bundle.xml /opt/jboss/server/default/deployed/nuxeo.ear/config
```

Then, go into the newly created folder, build the project and copy the new .jar:

```
mvn clean package
cp target/nuxeo-platform-login-portal-sso-VERSION.jar \
/opt/jboss/server/default/deploy/nuxeo.ear/system
```

Restart JBoss

33.2.2. Generate a sample project with nuxeo-archetype-portlet archetype

33.2.2.1. Install the archetype

Before creating a new project, you first need to install `nuxeo-archetype-portlet` on your local repository.

Checkout the sources:

```
svn co http://svn.nuxeo.org/nuxeo/org.nuxeo.archetypes/nuxeo-archetype-portlet/trunk \
nuxeo-archetype-portlet
```

Then, install the archetype on your local repository:

```
mvn clean install
```

33.2.2.2. Create a new project

To create a new project named `my-portlet` in the `com.company.sandbox` area:

```
mvn archetype:create -DartifactId=my-portlet -DgroupId=com.company.sandbox \
  -DarchetypeArtifactId=nuxeo-archetype-portlet \
  -DarchetypeGroupId=org.nuxeo.archetypes \
  -DarchetypeVersion=1.0-SNAPSHOT
```

There are two arguments:

- ArtifactId: usually the name of your project, with '-' to separate the words if there are many
- GroupId: the domain name of your project. Usually the package parent name of your classes.

Maven should have generated the following source layout:

```
   my-portlet
     |-- pom.xml
     `-- src
         `-- main
             |-- java
             |   `-- com
             |       `-- company
             |           `-- sandbox
             |               `-- portlet
             |                   `-- sample
             |                       `-- NuxeoSamplePortlet
             |-- resources
             |   `-- org
             |       `-- nuxeo
             |           `-- portlet
             |               `-- sample
             |                   `-- i18n
             |                       |-- SamplePortletMessages_en.properties
             |                       `-- SamplePortletMessages_fr.properties
             |
             `-- webapp
                 |-- index.jsp
                 `-- WEB-INF
                     |-- portlet.xml
                     |-- web.xml
                     |-- jsp
                     |   |-- edit.jsp
                     |   |-- sampleHelp.jsp
                     |   `-- sampleView.jsp
                     `-- tld
                         |-- c.tld
                         |-- fmt.tld
                         `-- fn.tld
```

33.2.3. Test the newly created portlet

Right now, you can build the portlet, but it won't deploy on Jahia. In fact, `portlet.xml` still references the wrong class.

Edit `portlet.xml` and find the line:

```
<portlet-class>org.nuxeo.portlet.sample.NuxeoSamplePortlet</portlet-class>
```

In our example, the class is now in the package `com.company.sandbox.portlet.sample`, change the line to reference the right class:

```
<portlet-class>com.company.sandbox.portlet.sample.NuxeoSamplePortlet</portlet-class>
```

You can now build the portlet and deploy it on Jahia. In your project folder:

```
mvn clean package
cp target/my-portlet.war /opt/jahia/tomcat/webapps/jahia/WEB-INF/var/new_webapps
```

Just wait a few seconds that Jahia finds your new portlet and deploys it. Then, go to Jahia and add the portlet. Go to the edit page of the portlet and fill the different informations (like Nuxeo Server URL, UserName, Password, ...).

When all is configured, write your name and send it. You should have "Hello your_name!" as response.

33.3. Developping Nuxeo Portlets

33.3.1. NuxeoPortlet class

Using NuxeoPortlet as base class make easier the developement of Portlets which have to communicate with Nuxeo EP.

NuxeoPortlet has some useful methods:

- error handling
- call a restlet on the configured NuxSeo server
- global preferences handling (informations stored in global preferences are shared with all the users of the portlet)

33.3.2. Project from nuxeo-archetype-portlet archetype

If you use the archetype to make your project, you will have some default behaviors:

- an administrator role is already defined (in web.xml and portlet.xml), so you can link, for instance, Jahia administrator role to your portlet administrator role. In your code, use the method `isAdministrator()` from NuxeoPortlet class to know if the current user is an administrator or not.
- a basic, but with all the informations needed to connect to a Nuxeo EP, edition page for the portlet is already done. All the informations are stored in global preferences for the portlet (through the `getGlobalPreferences()` and `saveGlobalPreferences()` methods). That means you can allow only administrators to modify these preferences, but they will be shared with all the users. You can of course

create your own edition page and override the default behaviour.

- the action requests are dispatched in convenient methods, depending of the portlet mode. For instance, in the class `NuxeoSamplePortlet`, we don't use the `processAction()` method, but `processViewAction()` when portlet is in VIEW mode and `processHelpAction()` when in HELP mode. You don't have to deal with the different portlet modes. Need to process a view action request? just use `processViewAction()`.

See the nuxeo-portlet-search project as an example of how to use or override the defaults behaviours. This portlet allows the user to make a simple search or an advanced one (like in Nuxeo EP) on the configured Nuxeo server.

33.3.3. portlet.xml

Beside the portlet class, `portlet.xml` lets you customize the portlet description, name, title.

There are also some default init parameters which are used to build the global preferences of the portlet. All the init parameters will be stored in the global preferences, so you can add the parameters you want and then use them in your portlet through the global preferences.

All the init parameters already in `portlet.xml` are needed by `NuxeoPortlet` to behave correctly.

33.3.4. Restlets

The portlet communicates with the Nuxeo server through some restlets. In our example, we call the restlet `sample` with a name as parameter.

To do a restlet call, use the method `doRestletCall()` from `NuxeoPortlet`, it makes your call and returns the result as a `Representation` (use `getText()` on it to have a `String` containing the result). The `doRestletCall()` method takes a `RestletCall` object as argument which contains the different parameters to do the restlet call: the restlet name, the path parameters and the query parameters.

See `NuxeoSamplePortlet` class or `nuxeo-portlet-search` project as example of how it works.

The `doRestletCall()` builds the URL to call from the configured Nuxeo server in the global preferences and from the parameters of the `RestletCall` object. In our example:

```
http://localhost:8080/nuxeo/sample?name=my_name
--------------------------- ------ ------------
             |                |         `-- the query parameters of the RestletCall object
             |                `-- the restlet name of the RestletCall object which will be
             |                    in the path parameters
             `-- the Nuxeo server URL from the global preferences
```

If all is well configured in the global preferences, you don't need to bother with authentication, just call the restlet.

To make a specific work, you have to develop your own restlets and contrib them to the `nuxeo-platform-restPack` project. Then rebuild the project, copy the .jar in your deployed Nuxeo EP and restart JBoss. You can now call your newly created restlets in your portlet.

33.4. Available portlets

33.4.1. Nuxeo Search Portlet

This portlet allows the user to search documents in a Nuxeo repository through a portlet container.

33.4.1.1. Installation

You just have to compile the portlet and install it into a portlet container, Jahia for instance, as explained in the previous sections.

After the installation of the prerequisites, you need to deploy on your portlet container the packaged portlet.

If you don't have already a packaged portlet, checkout the sources and package it:

```
svn co https://svn.nuxeo.org/nuxeo/nuxeo-addons/nuxeo-portlets/trunk/nuxeo-portlet-search \
nuxeo-portlet-search
cd nuxeo-portlet-search
mvn package
```

Then deploy the packaged portlet in Jahia, while Jahia is running:

```
cp target/nuxeo-portlet-search.war /opt/jahia/tomcat/webapps/jahia/WEB-INF/var/new_webapps
```

33.4.1.2. How to use Nuxeo Search Portlet

Before making any test, a Nuxeo EP should be running and you must configure the portlet for a Nuxeo EP:

- Attach a group or a user from Jahia (the one who will be considered as Administrator for the portlet) to the Administrator role of the portlet.
- Then, go to the edit view of the portlet and fill the different fields, so that the portlet can communicate with your Nuxeo EP.

This portlet provides the two search methods of Nuxeo EP, the simple one and the advanced one. They have the same behavior than the ones in Nuxeo EP.

Chapter 34. Desktop integration tools

This chapter is dedicated to modules to improve the user experience through integration of browser and office productivity software with a Nuxeo Enterprise Platform server.

The source code of the desktop integration tools is gathered here: `https://svn.nuxeo.org/nuxeo/desktop-integration`.

34.1. Drag and drop browser extensions

The login page of Nuxeo EP advertises two links to download and install browser extensions that makes it possible to import any files from the client's desktop file manager into the Nuxeo EP repository as new documents by drag and dropping the files or folders into the browser window.

Packaged installers for those extensions are made available on the `http://updates.nuxeo.org` URL.

34.1.1. Server side import service: the FileManagerService

All drag and drop desktop browser extensions use the `FileManagerService` and its contributed file importers to perform the actual document creation.

TODO: add here the list of remote interfaces to the FileManagerService API (REST, SOAP?).

34.1.2. Microsoft Internet Explorer plugin

The Internet Explorer plugin source code is available at: `https://svn.nuxeo.org/nuxeo/desktop-intergration/nuxeo-dragdrop-ie-extension` . This module is a coded using the C# language and needs the dotNET runtime version 2.0 or later.

TODO: give technical details on the WS protocol used to perform the file uploads.

34.1.3. Mozilla Firefox plugin

The Firefox plugin source code is available at: `https://svn.nuxeo.org/nuxeo/desktop-intergration/nuxeo-dragdrop-ff-extension` .

The firefox plugin is packaged as a regular XPI and is platform-independent and uses the built-in Javascript API of Mozilla to upload the content of the files as POST request to a RESTful web service URL.

TODO: give details on the remote RESTlets used for this implementation.

34.2. Online document editing with LiveEdit

LiveEdit is the generic name of a set of client- and server-side components meant to seemingly integrate remote Office document editing without having to rely of manual upload files through in-browser HTML forms.

34.2.1. Functional overview

The generic functional use case is the following: the users use a standard web browser to login and browse the web interface of Nuxeo workspaces. Upon a click on a link flagged 'edit online' on a Nuxeo document having an Office file as attachment, the user automatically gets the content of the file open for editing in the right client side application (e.g. Microsoft Word, OpenOffice.org, ...)

When saving changes, the new version of the file is automatically re-uploaded to the Nuxeo server through a SOAP or RESTful web service to update the original document content and make the changes available to the

other users of the workspace.

The version number can be incremented upon LiveEdit editing. A lock can be optionally set on the original document so that two users do not overwrite each other changes in concurrent LiveEdit sessions.

LiveEdit components should also support 2 auxiliary use cases:

- creation of a new Nuxeo document with an empty client-side generated attachment of the specified mimetype;
- creation of a new Nuxeo document with an client-side generated attachment preinitialized with the copy of an existing binary attachment stored elsewhere.

The following introduce the details of the 3 use cases implemented by the LiveEdit system along with the technical components at work.

34.2.2. Functional use cases

34.2.2.1. Editing the attachment of an existing document of the Nuxeo repository

The user wants to a edit a non-empty Office file stored as a property of a Nuxeo document.

The user authenticate to Nuxeo (login/password or some implementation of SSO) with her web browser (IE or Firefox) and browse her workspaces till the summary view of a document she has the rights to edit.

If the document has Blob storing property containing a non-empty office file with a mimetype flagged as live editable (MS Office and OpenOffice related mimetypes), the web interface generate a link (marked 'edit online') that automatically opens the right desktop application with the content of the attachment in opened in the editing window.

Upon saving, the desktop editor opens popup leaving the user with the option either to save a local copy or the save on the Nuxeo server. For the latter the user can also choose to increment the minor or major version number or to overwrite the current version.

Upon completion the user can check by browsing back to her document that attachment has be updated with her changes on the Nuxeo server.

34.2.2.2. Creating a new document with an empty attachment

The user wants to create a new Nuxeo document from scratch directly from an Office productivity application.

The user authenticate to Nuxeo (login/password or some implementation of SSO) with her web browser (IE or Firefox) and has the possibility to click on a document creation menu with the following items:

- New Word document
- New Powerpoint presentation
- New Excel spreadsheet
- etc. (similar options for OpenOffice.org apps)

Clicking one of those links automatically opens the right desktop application with a new empty document opened in the editing window.

Upon saving, the desktop editor opens popup leaving the user with the option either to save a local copy or the save on the Nuxeo server. For the latter the user has to choose among a flat list of candidate remote locations selected are labeled with both a title (e.g. 'My Workspace') and a path to the location (e.g. `/default-domain/workspace/my-workspace`)

Upon completion the user can check by browsing to her new document at the selected location. The file attachment of that document has the content of the saved Office file.

34.2.2.3. Creating a new document preinitialized with an existing attachment

The user wants to create a new Nuxeo document directly from an Office productivity application but reusing the content of an attachment stored as a property of an existing Nuxeo document.

The user authenticate to Nuxeo (login/password or some implementation of SSO) with her web browser (IE or Firefox) and browse her workspaces till reaching a document with a non-empty Office file attachment. The view of that file carries a link labeled `'edit online as a new document'`.

Clicking on that link opens the right desktop application preinitialized with a copy of the content of the template document attachment.

Upon saving, the desktop editor opens popup leaving the user with the option either to save a local copy or the save on the Nuxeo server. For the latter the user has to choose among a flat list of candidate remote locations selected are labeled with both a title (e.g. `'My Workspace'`) and a path to the location (e.g. `/default-domain/workspace/my-workspace`)

Upon completion the user can check by browsing to her new document at the selected location. The file attachment of that document has the content of the saved Office file.

34.2.3. Architectural overview

34.2.3.1. Overview of the components

The LiveEdit feature is built upon a set of the following interacting components:

- webapp components to generate a link to a generated XML bootstrap file describing the file to edit remotely along with connection parameters. Such a sample bootstrap file is provided in the technical annexes of this specification file.
- a browser protocol handler (i.e. plugin for Internet Explorer or Firefox) that parses the xml bootstrap size and launch the relevant desktop application according to the mimetype
- an editor plugin for each desktop application (MS Office, OpenOffice) to be able to make the desktop application fetch the file from Nuxeo through SOAP or REST GET with connection parameters provided in the bootstrap file and save it back to the server using SOAP or REST POST as well.

  Before and after editing the document, the editor plugin can fetches a list of configurable actions to call on the server through SOAP or REST GET (lock/unlock, check in/ check out, validate, etc.).
- a web service component (EJB3 stateful bean with JAXWS extensions) to implement the required methods along with the WSDL definition need by the desktop client

The source code of the various LiveEdit client side components is available at:
`https://svn.nuxeo.org/nuxeo/desktop-intergration/nuxeo-liveedit-*` .

34.2.3.2. The Bootstrap client module (part 1)

The Bootstrap module must intercept the click on the "online edit" link using a dedicated protocol handler packaged as a browser plugin.

The "online edit" link has the following pattern:

```
nxedit:http://localhost:8080/nuxeo/nxliveedit.faces?action=edit&repoID=[repoID]&docRef=[docRef]&schema=
```

The protocol handler will be called by the OS/Browser and receive the URL. In turn it will receive the XML bootstrap file.

In case of create use cases the previous patterns change as follows:

- user case #2: no docid but need to provide the type of the future document and field location of the blob to be stored in:
  - `nxedit:http://localhost:8080/nuxeo/nxliveedit.faces?action=create&repoID=[repoID]&mimetype=[mime`
- user case#3: docid and field path of the original blob AND doctype and fieldpath to of the document that will host the result:
  - `nxedit:http://localhost:8080/nuxeo/nxliveedit.faces?action=createFromTemplate&templateRepoID=[te`

The Bootstrap client module rewrites each of those URIs as valid HTTP GET by swapping the prefix:

- `http://localhost:8080/nuxeo/nxliveedit.faces?[query_parameters]`

The Bootstrap client protocol must strip the "nxedit:" prefix of the URI to get the HTTP URL and thus work either with SSL or not:

- `nxedit:http://localhost:8080/nuxeo/nxliveedit.faces?[query_parameters]`

should be transformed into:

- `http://localhost:8080/nuxeo/nxliveedit.faces?[query_parameters]`

while:

- `nxedit:https://localhost:8080/nuxeo/nxliveedit.faces?[query_parameters]`

should be transformed into:

- `https://localhost:8080/nuxeo/nxliveedit.faces?[query_parameters]`

In order to generate valid `nxedit:` URIs it is strongly advised to use the JSF functions defined in the `org.nuxeo.ecm.platform.ui.web.tag.fn.DocumentModelFunctions` class.

The functions are available under the `nxd` (`http://nuxeo.org/nxweb/document`) namespace:

- `liveEditUrl(DocumentModel)` : get the `nxedit:` URL to edit a document file attachment (default File-like types)
- `liveEditUrl(DocumentModel, String, String, String)` : get the nxedit: URI to edit a document providing schema, blob field and filename field names
- `liveCreateUrl(String)` : get the nxedit: URI to create a new document of type File providing the mimetype as argument
- `liveCreateUrl(String, String, String, String, String)` : get the nxedit: URI to create a new document with parameters: mimetype, doctype, schema, blob and filename field names
- `liveCreateFromTemplateUrl(DocumentModel)` : get the nxedit: URI to create a new document of type File reusing the content of the blob of the provided template DocumentModel (assumed to have the "file" schema).
- `liveCreateFromTemplateUrl(DocumentModel, String, String, String, String, String, String)` : get the nxedit: URI to create a new document from template. Parameters are: template DocumentModel, template schema, template blob field, target document type, target schema, target, blob field name, target filename field.

34.2.3.3. The Bootstrap server module

The Bootstrap server module will be a simple Seam component called using the info passed as request parameters to generate the XML payload of the bootstrap file.

The boostrap server module is currently located in webapp-core here:

- `org.nuxeo.ecm.webapp.liveedit.LiveEditBootstrapHelper`

34.2.3.4. Bootstrap data download

The HTTP response to that GET URL is a bootstrap file containing an XML payload. This file contains:

- the action ID so the client knows to interpret it
- repo name
- document unique reference (in case of editing) or template reference (new from template)
- the document Nuxeo type of the document to create or edit
- the fieldpath hosting the binary attachment
- associated filename
- associated mime-type
- principal name
- whether the result of this editing session should be saved as a new Nuxeo document else where in the repository (creation use cases)

In case of creating from a sample document:

- the document id of the template
- the fieldpath hosting the binary attachment to initialized the editor with

The XML payload further contain a copy of all the HTTP request headers and cookies + basic auth credentials and the adress of the WSDL description of the LiveEdit webservice.

Please refer to the sample XML bootstrap file in the annexes for more details on the syntax. Some fields (eg. document reference) might be empty or missing in case of document creation (use cases #2 and #3).

34.2.3.5. The bootstrap client module (part 2)

The bootstrap module receives and parses the content of the XML bootstrap file. According to a set of configurable rules the bootstrap module launch the right editor with bootstrap file as command line parameter.

34.2.3.6. Authentication management during bootstrap

The Bootstrap client will need to do an http call to get the xml file from the server. This call must be authenticated. So the protocol handler must reuse the browser session.

34.2.3.7. The client editor and its plugin

In case of document editing (use case #1):

- call WS to get list of pre-edit actions
- display a dialog for letting user select action
- call WS to download the file

- call WS to get list of post-edit actions
- display a dialog for letting user select action
- save and upload the file to Nuxeo Server
- terminate (close the WS session)

In case of document creation (use case #2 and #3):

- call WS to get list of pre-edit actions
- display a dialog for letting user select action
- call WS to download the template file (use case #3)
- call WS to get list of post-edit actions (e.g. choose title)
- call WS for the list of candidate server locations
- display a dialog with actions and dropdown list of candidate locations
- create new document and upload the file to Nuxeo Server
- terminate (close the WS session)

34.2.3.8. Authentication of the client editor

All WS requests (SOAP and RESTful) from the editor plugin back to the WS server should reuse all the HTTP cookies along with any basic auth parameters to ensure the request will pass through any authenticating reverse proxies (e.g. CAS, mod_sso, ...) as if they were the original browser.

34.2.4. The Web Service component

It is responsible to answer the WS calls of the editor client. Most of its business logics should be defined has an overridable extension point so that customer project can change most of the LiveEdit global behavior without having to re-compile / re-package the client part.

In particular the list of candidate locations to 'save as new document' is provided by the WS server-side API to the LiveEdit client. The list should default to the list of Workspaces the user currently has the "AddChildren" permission. The WS server should dynamically compute that list according to an extensible service (i.e. overridable by a extension point) so that customer project can register a custom Java class that is responsible to implement the custom business logics in case the list of workspaces is not enough.

The location selected by default should also be defined on the server side and overridable by the same extension point.

34.2.5. More on editor launch

Based on a configuration file containing mimetype/editor mapping, the bootstrap module will launch an editor. This configuration file should look like that:

```
<editors> <editor name="MSOOfficePlugin">
<pluginType>.net<pluginType> <mime-types>
<mime-type>application/msword</mime-type> ...
</mime-types> </editor> <editor
name="OOfficePlugin">
<pluginType>exec<pluginType> <mime-types>
<mime-type>application/vnd.oasis.opendocument</mime-type>
... </mime-types> </editor> ... </editors>
```

This is very important that bootstrap client can be separated from the editors plugins, because there will plugins contributed for specific editors. The simplest and most neutral way of launching an editor plugin is just executing the editor plugin passing it a copy of the bootstrap file. This file will be the same as the one returned by Nuxeo server with additional authentication information: cookies and Login/Password.

34.2.6. More on pre- and post-editing actions

Actions available on the document may depend on the custom project specifications, and it is important that it is totally transparent for the client plugin UI: we don't want to build a version of the client plugins for each project.

So the idea is that the WebService will provide the client editor plugin a simple list of actions, the client will simply display available actions, and eventually ask the server to execute them without knowing the underlying logic.

This "action logic" is somehow close to what we already do in the web layer, an action defines an action id and a label.

If several actions can be done (like checkout + lock), then they will be combined as compound actions: checkout, lock, checkout_and_lock. This way we won't have to handle associations conditions on the client side: the user can always select at most one action: none / lock / checkout / checkout_and_lock

34.3. Configuring LiveEdit links

34.3.1. Configuration policies

Starting with Nuxeo 5.1.6, the liveEdit link display policy is pluggable.

This means you can define on what kind of documents the liveEdit links will be displayed.

Nuxeo supports 3 different policies :

- Client based configuration

  The client browser tells the server what kind of documents can be "live edited".

- Server based configuration

  Links are only displayed on mime-types that are configured on the server to be live editable.

- Mixed configuration

  Links are only when server and client policy both apply.

34.3.1.1. Client based configuration policy

For that the firefox and msie plugins add the live editable mime-types to the accept header sent by the browser. This way, the server can decide on what type of file the liveedit link must be displayed. If you use Firefox, just upgrade to the last version and use the configuration pannel to tell on what mime-types you want liveedit links.

The main advantage of this policy is that only the users that have LiveEdit plugin installed will see the liveEdit links.

The accept header send by a LiveEdit enabled client looks like that :

```
...,application/x-nuxeo-liveedit;application/vnd.oasis.opendocument.text;application/vnd.oasis.opendocument.pres
```

34.3.1.2. Server based configuration policy

In this case the server will only display the LiveEdit links for files that have a mime-types that is declared "LiveEditable" in the MimeTypesRegistry service.

34.3.1.3. Mixed configuration policy

Simply display links on mimetypes that are livededitable on the client side and on the server.

34.3.2. Changing the configuration policy

Default policy is client based, in order to change that, the only thing you have to do is edit the nuxeo.properties file in nuxeo.ear/config.

```
# LiveEdit configuration detection (client/server/both)
org.nuxeo.ecm.platform.liveedit.config=client
```

Chapter 35. Nuxeo WebDAV interface

The `nuxeo-platform-webdav-server` module provides a WebDAV interface on top of Nuxeo services.

35.1. WebDAV clients

There are several WebDAV clients available. Each of them has its specific behavior, and also some limitations.

35.1.1. Path vs displayName

In WebDAV specification each accessible resource must have a name (part of the URL) and can have a displayName (defines how the resource should be displayed to the user). Unfortunately, most WebDAV client don't use the displayName property even if they ask the server to send it. Because of this limitations, most client use the last part of the URL as display name (with all limitations due to URL encoding). In addition of the display problem, this has side effect on other features like renaming. When trying to rename a resource (for example after creating a new Collection/Folder), most clients will send a move operation (change the path) instead of sending a propatch on displayName.

Inside Nuxeo 5, the path (id) and the display name (title) are clearly separated:

- id/path is generated from the title but is not equal: non standard characters are escaped and length is limited to avoid having too long URLs.

  Foe example, a folder with title "New Folder" will have "new-folder" as a name.

- The name has to be unique within a given container, but 2 documents can have the same title.

35.1.2. Filesystem resource vs Nuxeo DocumentModel artifact

Inside WebDAV, each container is a collection. Most WebDAV client consider that a collection is not a document, but a simple folder. This limitation is important because Nuxeo supports folderish documents, and document types that can contain several files. Because of that, depending on the document types we may want to map, a Nuxeo document could be seen as a collection or as a resource. Unfortunately, since all clients consider collections as simple folders, this is not very useful.

35.1.3. MS Web Folder client

One of the most popular WebDAV client are Web Folders that are included inside Windows OS. Unfortunately depending on the version of Windows (including service pack) and also on the version of MSOffice the webdav client will be implemented in a different way with different bugs and behavior. Web Folders are not the worse WebDAV client, but they are definitely the less predictable...

For a complete bug listing of Web Folders depending on version please see
http://www.greenbytes.de/tech/webdav/webdavfaq.html,
http://www.greenbytes.de/tech/webdav/webfolder-client-list.html,
http://www.greenbytes.de/tech/webdav/webdav-redirector-list.html

35.2. Fooling WebDAV clients

Because of the path vs name problem, the Nuxeo 5 WebDAV connector introduce some hacks to force some webdav client displaying the correct names.

This hacks can be activated of not depending on the User-Agent header provided by each client.

35.2.1. Available hacks

- Virtual path for lief resources (ie: non collections)

  Because most client uses the last part of the URL to get the displayName, a simple hack to have resources displayed correctly is to have a specific URL for them.

  By default a resource with URL http://nxServer/nuxeo/dav/default/default-domain/workspaces/my-workspace-1/file-1 will be displayed as "file-1". But if file-1 contains a file named "MyFile.doc", when using a virtual URL like http://nxServer/nuxeo/dav/default/default-domain/workspaces/my-workspace-1/file-1/WebFolder_FileName/MyFile. then the display will be correct.

- Use filename as resource name

  The resource URL will be http://nxServer/nuxeo/dav/default/default-domain/workspaces/my-workspace-1/file-1/MyFile.doc. This is very much like the previous hack, but this is harder to resolve on the server side, because there may be several file resources containing MyFile.doc in the same container.

- Use fake get parameters for collections names

  Getting the good displayName for collection resources is a little bit harder, previous hacks won't work.

  In some webdav client, the url parsing is so weak, that we can use URLs like : http://nxServer/nuxeo/dav/default/default-domain/workspaces/my-workspace-1?displayName=/My%20Workspace% to have the collection correctly displayed as "My Workspace 1". Although this hack does not work with most clients, it works with most Web Folders version.

35.2.2. Configuring Nuxeo WebDAV connector for each client.

The Nuxeo WebDAV connector enables you to configure for each WebDAV clients:

- The hacks you want to enable
- If you want full or relative URLs
- If your client needs MS specific DAV Headers (required for some versions of Web Folders)

Configuration can be contributed using a simple Extension Point.

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.platform.webdav.config.defaultContrib">

  <extension target="org.nuxeo.ecm.platform.webdav.config.WebDavConfigurationService" point="DavClientConfigurat
    <documentation>
      Configuration for MS Web Folders (both versions) and WebDrive.
    </documentation>
    <davClientConfiguration name="MSWebFolders" enabled="true">
      <needGetParameterForCollectionNamming>true</needGetParameterForCollectionNamming>
      <needVirtualPathForLief>false</needVirtualPathForLief>
      <useFileNameAsRessourceName>true</useFileNameAsRessourceName>
      <needFullURLs>true</needFullURLs>
      <needMSDavHeader>true</needMSDavHeader>
      <userAgentPatterns>
        <pattern>Microsoft-WebDAV</pattern>
        <pattern>Microsoft Data Access Internet Publishing Provider</pattern>
      </userAgentPatterns>
    </davClientConfiguration>
   ...
```

As you can see, each configuration set is attached to one or more User-Agent substring. Unfortunately, Web Folders clients do not send information about their version, even if it would be very helpful.

35.3. Nuxeo EP WebDAV implementation

35.3.1. Nuxeo EP WebDAV-specific features

Because of the limitations explained earlier, the Nuxeo EP WebDAV implementation provides some features to be able to work with existing WebDAV clients.

35.3.1.1. VirtualPath management

When URLs hack are activated, the clients will be able to display the resources with their displayName (title) instead of the name. But some of them will be fooled enough to use this displayName in URLs to do propfind or move calls. Because of that, the server must be able to resolve URLs that are constitued partially of path and partially of displayNames like:
http://nxServer/nuxeo/dav/default/default-domain/workspaces/My%20Workspace%201.

The Nuxeo EP WebDAV connector implements a custom URL resolver that is able to resolve these URLs. In also maintains a cache of these virtual URLs to speed up resolution (and also make it more consistent is case of name collisions).

35.3.1.2. WebDAV resource adapters

Displaying your Nuxeo documents as simple files (that's what will do most WebDAV client) can be very restrictive.

The Nuxeo WebDAV connector uses DocumentModel adapters to define how a Nuxeo Document must be mapped to a WebDAV resources. This adapter defines:

- How the display name of your resource will be generated.

  Simple document title or filename of the main file field...

  The default built-in behavior will be to return the filename for each document that has the file schema and otherwise return the title.

- How your DocumentModel will respond to a GET request.

  The default built-in behavior will be to return:

  - a folder listing as HTML for each folderish resource
  - the binary file for Documents that uses the file schema
  - the XML export for other non folderish resources

New DavAdapters can be contributed to define specific WebDAV mapping for your document types. For that, the Nuxeo WebDAV connector provides an extension point to let you register a new class implementing`org.nuxeo.ecm.platform.webdav.adapters.DavResourceAdapter` and associate it to a document type.

```
<extension target="org.nuxeo.ecm.platform.webdav.config.WebDavConfigurationService"
      point="DavAdapter">

  <davAdapter name="NoteDavAdapter" enabled="true">
    <typeName>Note</typeName>
    class>org.nuxeo.ecm.platform.webdav.adapters.NoteDavResourceAdapter</class>
  </davAdapter>
</extension>
```

35.3.2. Known limitations

Locking management is restricted to exclusive write locks, exactly as in Nuxeo Core.

Propset is for now not implemented as none of the used WebDAV client seems to use this method.

WebDAV versioning extension are not implemented.

35.4. Using the Nuxeo WebDAV connector

35.4.1. Installing the WebDAV connector

The Nuxeo WebDAV Connector is an optional additional component for Nuxeo EP.

In order to install the WebDAV features, you need to download the WebDAV Connector jar (or build it from source), copy the jar in `nuxeo.ear/system` and restart your server.

The WebDAV connector works for both 5.1 and 5.2 versions of Nuxeo EP, you just need to download the associated version or build it using the right POM file.

35.4.2. Connecting a client to Nuxeo WebDAV connector

The WebDAV URL for the default domain is http://$NuxeoServer/nuxeo/dav/default/default-domain.

In order to use MS Web Folders, you just need to go to "My Network Places" and choose "Add a new Network place". You can then enter the Nuxeo WebDAV URL and login/password.

Chapter 36. Reporting: Eclipse BIRT Driver

36.1. Overview

BIRT is an open source Eclipse-based reporting system that integrates with Java/J2EE application to produce compelling reports. The BIRT driver for Nuxeo enables BIRT to be used as reporting engine for the Nuxeo Content Repository. It basically gives an easy way to query the repository and create report from the results. Thanks to BIRT, reports can be run inside Eclipse or as servlets on the server side.

36.2. How to use it

In a Eclipse instance with Business Intelligence and Reporting Tools installed (http://www.eclipse.org/birt/phoenix/), deploy the following Nuxeo plugins:

- org.nuxeo.common
- org.nuxeo.ecm.client
- org.nuxeo.ecm.core.api
- org.nuxeo.ecm.core.query
- org.nuxeo.ecm.core.schema
- org.nuxeo.ecm.jboss_connector
- org.nuxeo.ecm.platform.search.api
- org.nuxeo.ecm.platform.usermanager.api
- org.nuxeo.logging
- org.nuxeo.runtime
- org.nuxeo.runtime.config
- org.nuxeo.birt.oda.nuxeoep
- org.nuxeo.birt.oda.nxueoep.ui

When eclipse is restarted a new "Data Source Should be available"

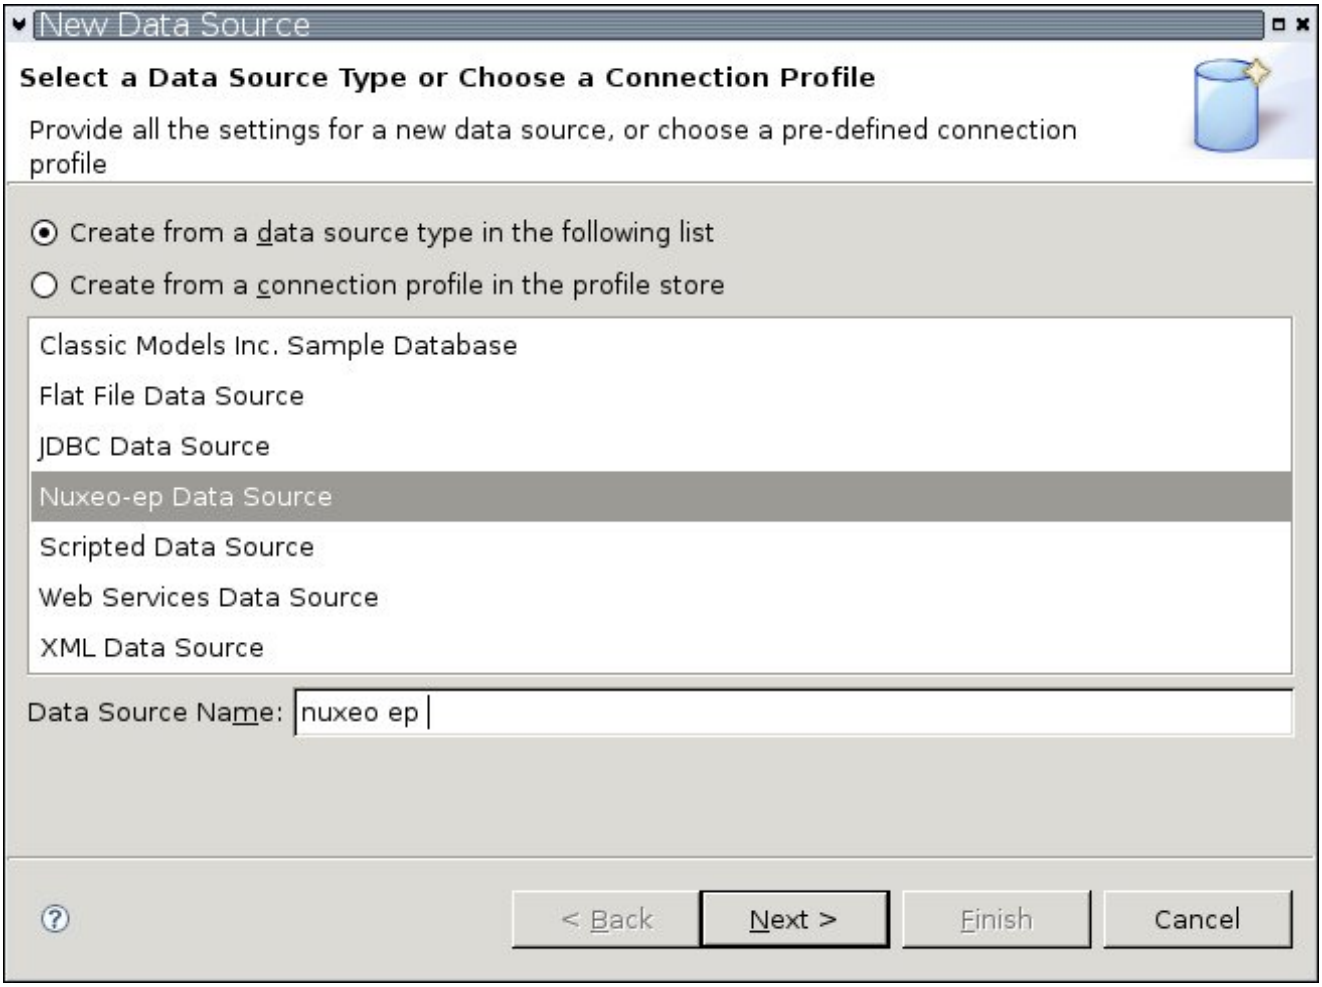
New Data Source
Select a Data Source Type or Choose a Connection Profile
Provide all the settings for a new data source, or choose a pre-defined connection profile
Create from a data source type in the following list
Create from a connection profile in the profile store
Classic Models Inc. Sample Database
Flat File Data Source
JDBC Data Source
Nuxeo-ep Data Source
Scripted Data Source
Web Services Data Source
XML Data Source
Data Source Name:
nuxeo ep
< Back
Next >
Finish
Cancel

Figure 36.1. In the Data Source Type screen, select "Nuxeo Data Source":

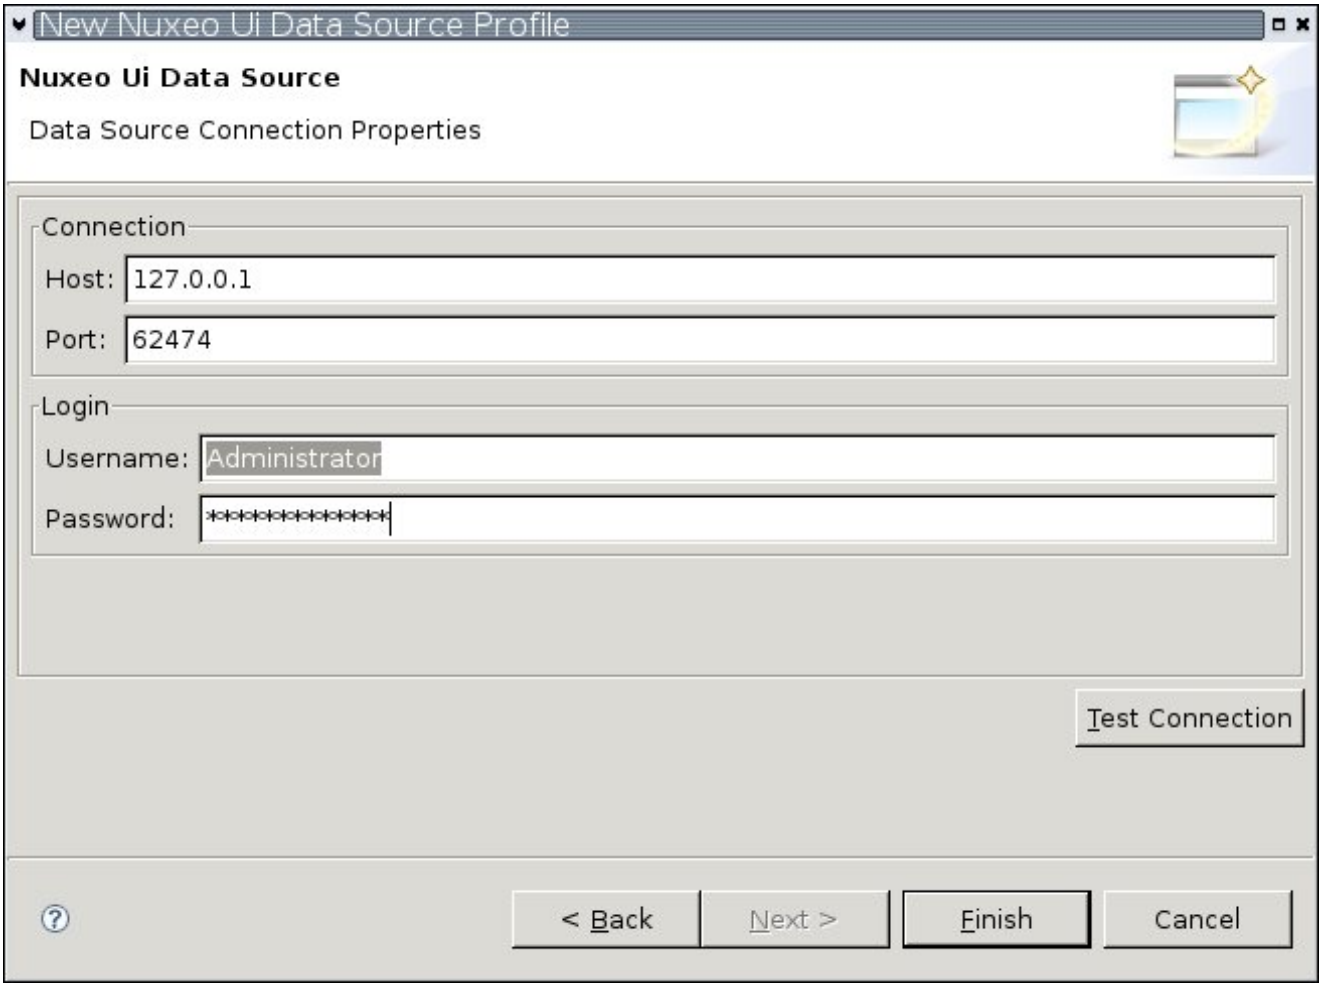

Figure 36.2. Fill login information

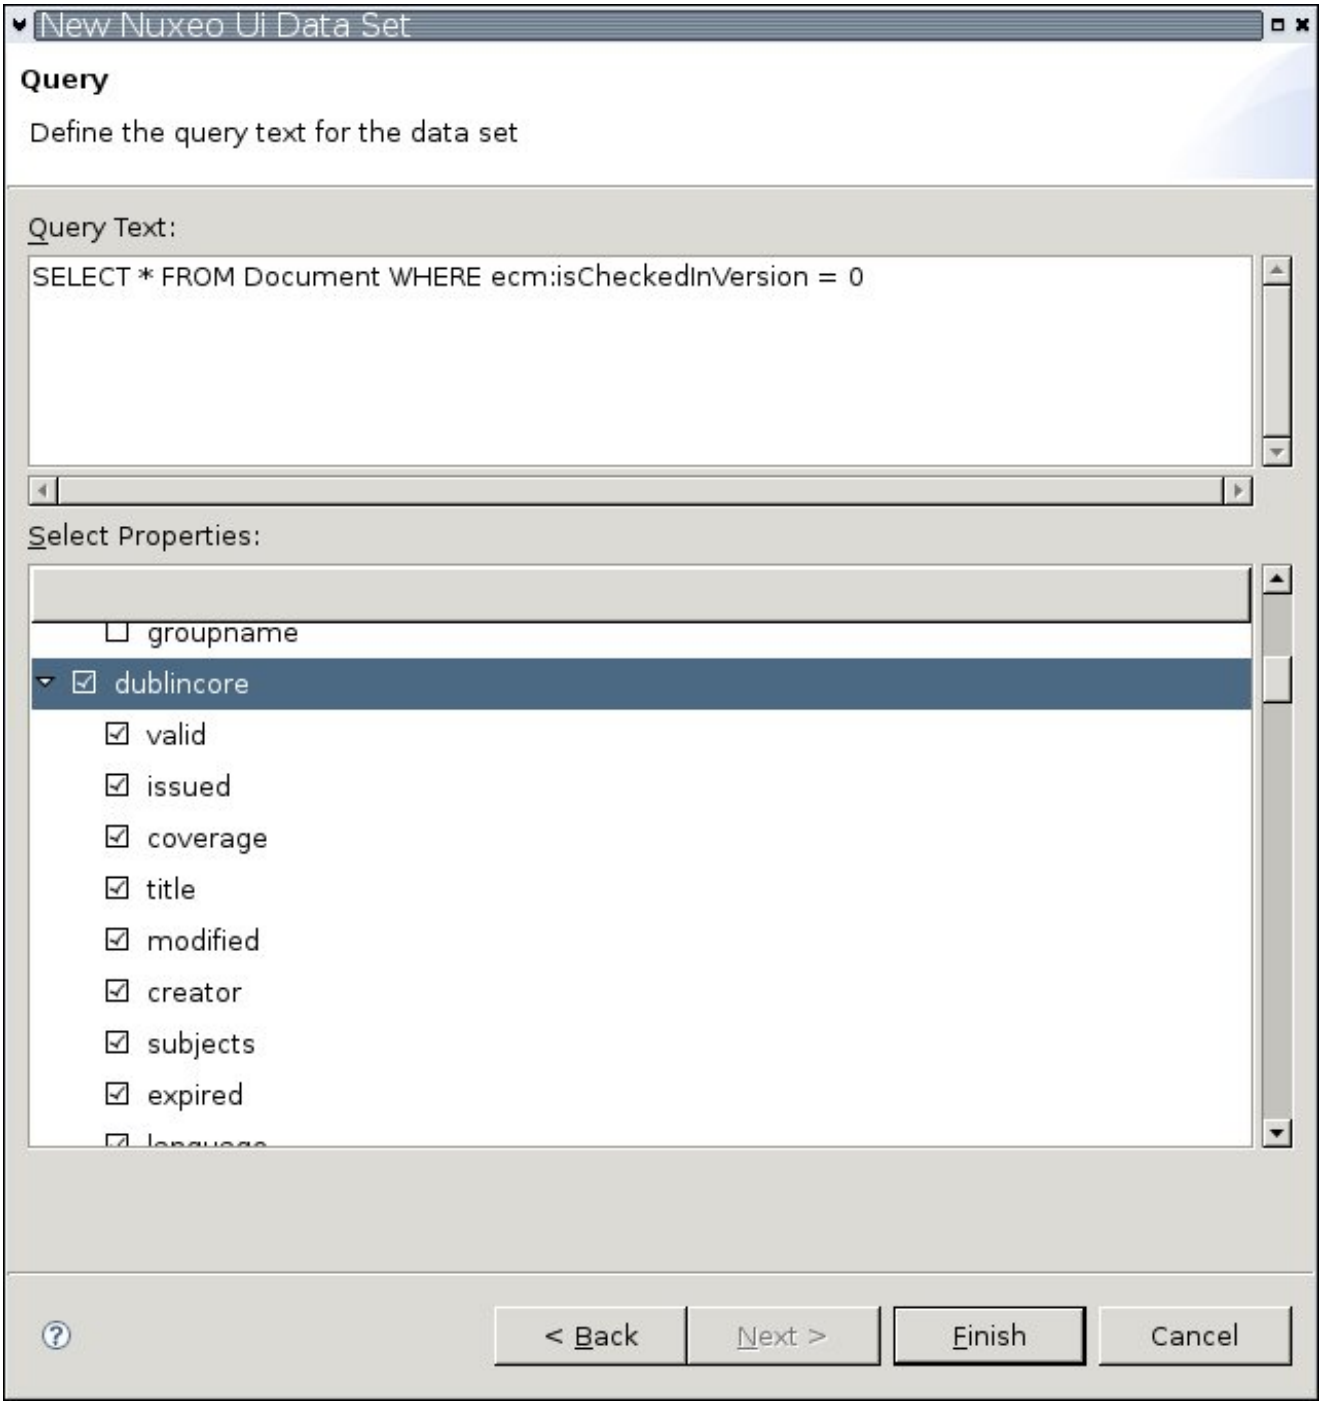

Figure 36.3. In the Data Set dialog, type NXQL query and select fields & schemas you would like to use in the report

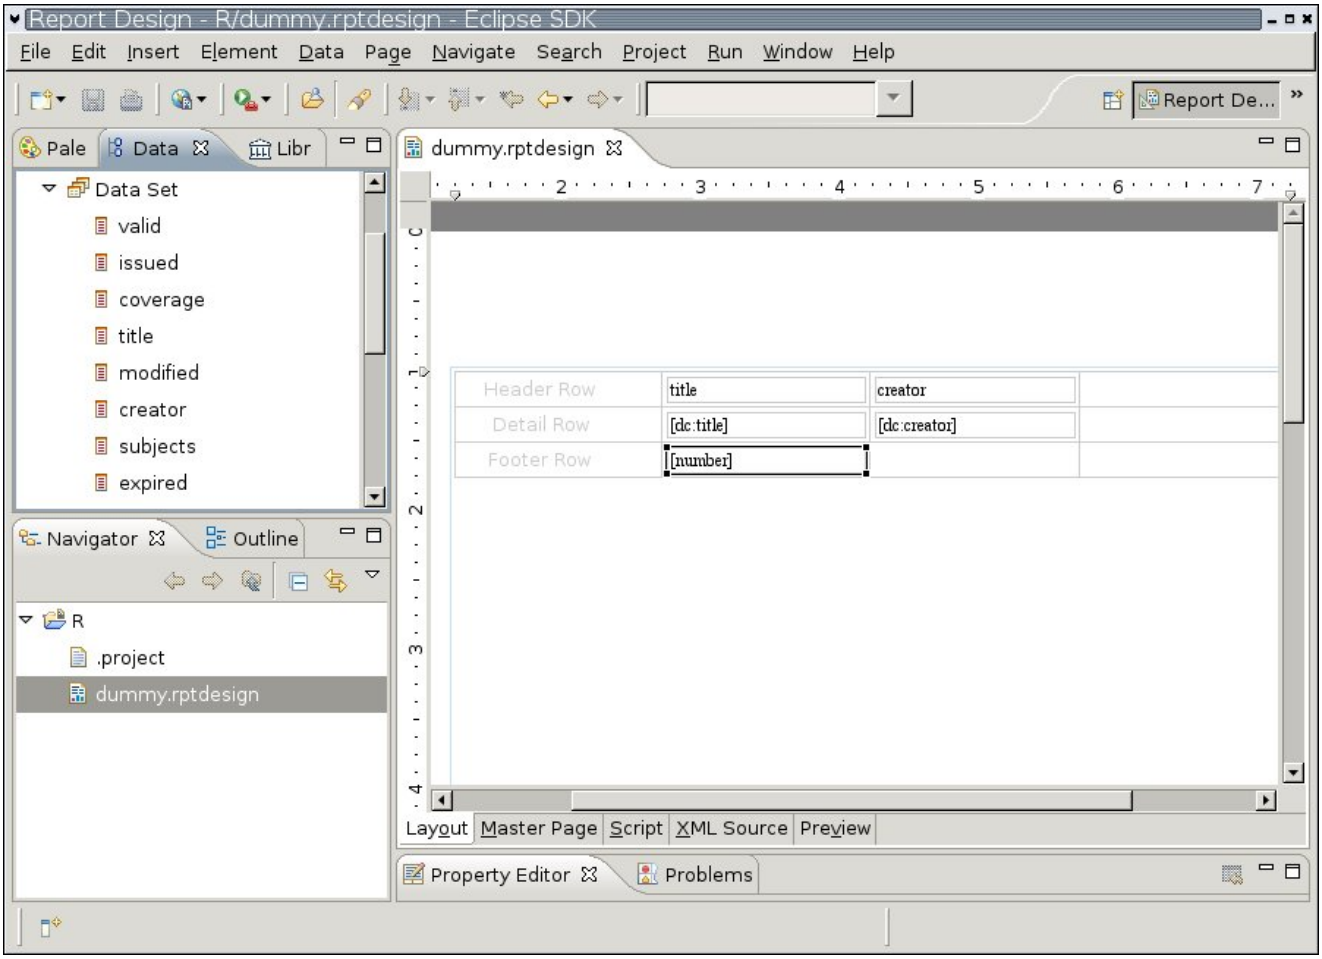

Figure 36.4. Data set is ready to be used in report

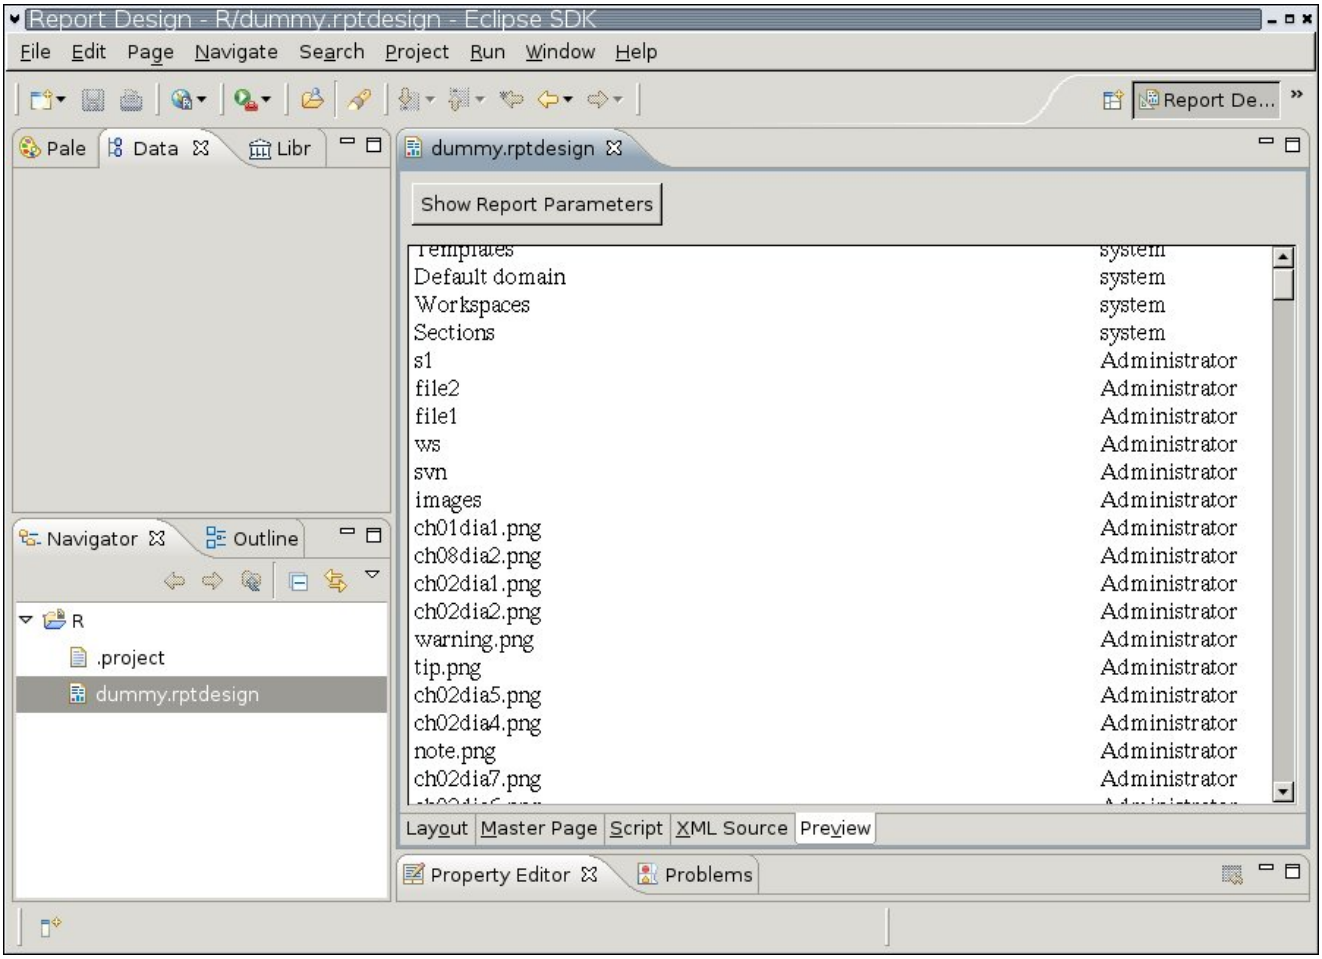

Figure 36.5.

36.3. Tomcat integration HOWTO

The new created report can be deployed in Tomcat to be available online:

1. Install the birt-viewer application following instruction from http://www.eclipse.org/birt/phoenix/deploy/viewerSetup.php

2. Deploy the plugins listed in previous section (except org.nuxeo.birt.oda.nxueoep.ui) in $TOMCAT_HOME/webapps/birt-viewer/WEB-INF/platform/plugins

3. Just copy the report file from workspace to $TOMCAT_HOME/webapps/birt-viewer/report

4. The report is available with an URL similar to http://localhost:13000/birt-viewer/frameset?__report=report/dummy.rptdesign

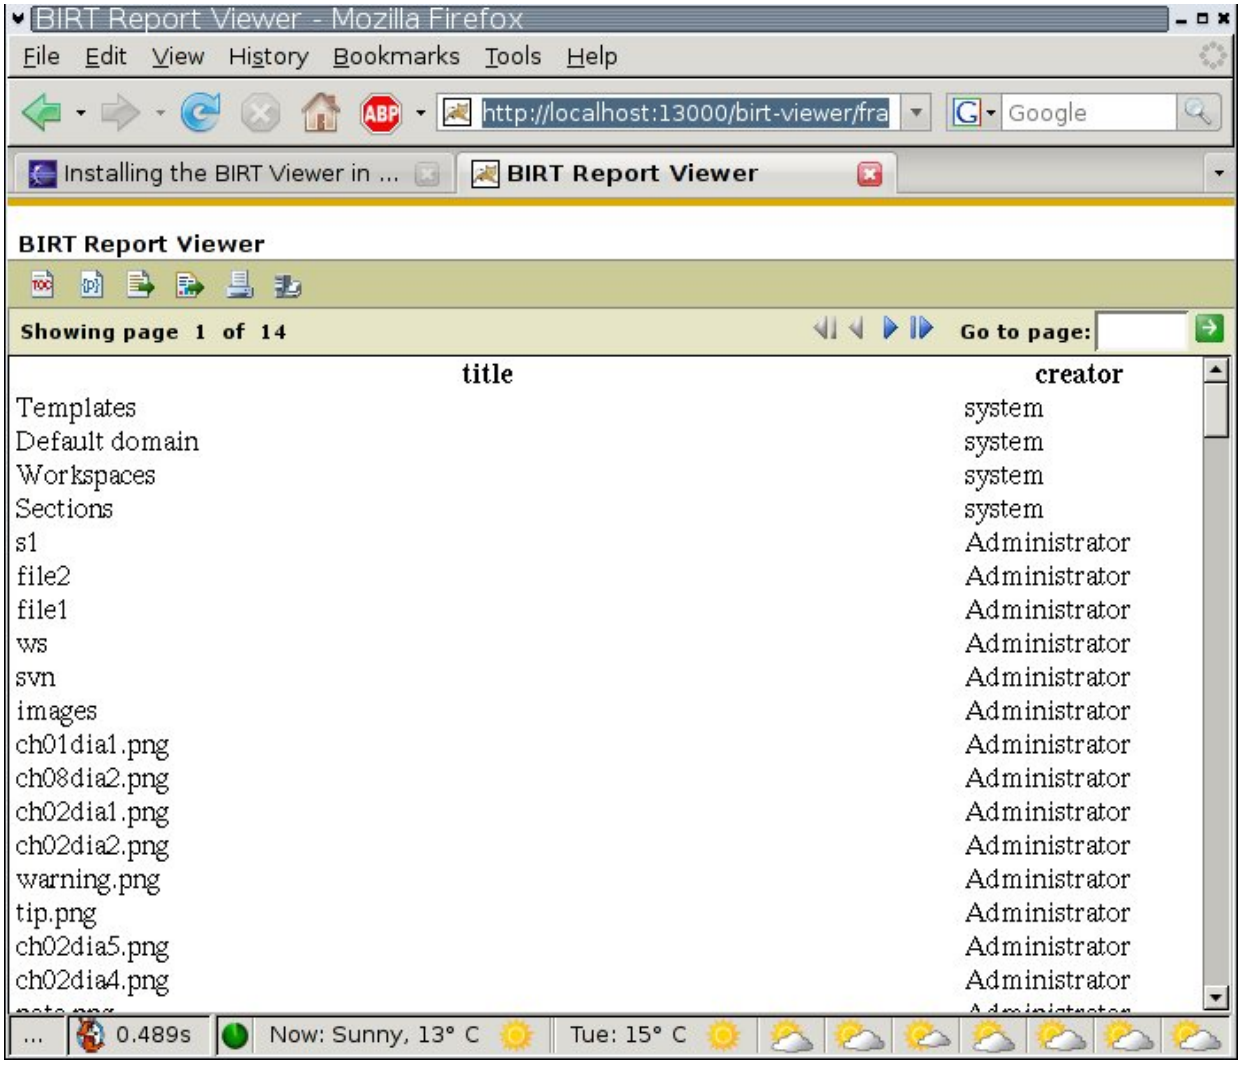
BIRT Report Viewer - Mozilla Firefox
File Edit View History Bookmarks Tools Help
http://localhost:13000/birt-viewer/fra
Google
Installing the BIRT Viewer in ...
BIRT Report Viewer
BIRT Report Viewer
Showing page 1 of 14
Go to page:
title
creator
Templates system
Default domain system
Workspaces system
Sections system
s1 Administrator
file2 Administrator
file1 Administrator
ws Administrator
svn Administrator
images Administrator
ch01dia1.png Administrator
ch08dia2.png Administrator
ch02dia1.png Administrator
ch02dia2.png Administrator
warning.png Administrator
tip.png Administrator
ch02dia5.png Administrator
ch02dia4.png Administrator
0.489s
Now: Sunny, 13° C
Tue: 15° C

Figure 36.6.

Chapter 37. Nuxeo Flex Connector

The `nuxeo-platform-flex` addon provides a Flex Connector to use Nuxeo services and Nuxeo Seam components as Flex Remote Services.

37.1. Overview

The Nuxeo Flex Connector uses Granite Data Services wich is is a free, open source (LGPL), alternative to Adobe® LiveCycle® (Flex™ 2+) Data Services. GraniteDS provide a full support of AMF3/Remote Object for EJB3/Seam/Spring/Guice/Pojo technology. Since it is highly configurable, we added support for Nuxeo Runtime Services.

37.2. Development environment

You will need Flex and Air sdk to build the flex connector and its applications You can download those right here:

- Flex SDK
- Air SDK

Check that your environment variable AIR_HOME and FLEX_HOME are setup correctly since Maven uses them to build your project.

Like all nuxeo project, you will need mercurial, maven and ant.

37.3. Build and Deploy

First, you have to get the sources on Mercurial.

```
hg clone http://hg.nuxeo.org/addons/nuxeo-platform-flex
```

Next step is building those sources with maven and deploy it on your server. Configure your build.properties file so ant knows your server's location.

```
ant install
```

This will run `mvn install` to build the project and ant copy to deploy the packages on your server.

By now you should have your connector deployed. Let's jump in the `sample` directory where you will find different simple examples. Run `ant install` to deploy them on your server. Next is a small overview of those samples. Start JBoss and go to http://localhost:8080/nuxeo/ to test them.

37.3.1. Sample Overview

- Actions sample Take a look at all actions register for the selected Document.
- Browser sample Browse through nuxeo Document and take a look a their properties.
- DocumentAPI sample Modify a Document's property.
- Facet Explorer sample Calls the SchemaManager Service to get Document Types implementing the given facet. Here to show the use of Nuxeo Runtime Service. This is a list of usual facet to test this sample: Folderish, Versionable, Orderable, Downloadable, Publishable, HiddenInNavigation,

BrowseViaSearch.

- Flex-login sample Log in Nuxeo Server and see the list of availables samples.
- Simple sample Ping the server and play with Document properties
- Tree sample Browse through Nuxeo documents with a simple navigation tree.
- Users sample Search and Modify a User
- Vocabularies sample Navigate through Nuxeo's Vocabularies.

37.4. Dive In

In this chapter, we'll see how to configure and call data services from Flex.

37.4.1. Data Services Configuration

Granite DS requires services configuration on both client and server side. To avoid painfull duplication of code, the connector provide two default factories and a default channel. If you are use to GraniteDS, you might wonder how the server side configuration is loaded. This is explained in Granite Configuration chapter.

Default Nuxeo Factory

There are two Different Factory, one to get Seam component and one to get Nuxeo services.

```
     <factories>
         <factory id="seamFactory" class="org.nuxeo.ecm.platform.ui.granite.factory.NuxeoSeamServiceFactory" / >
         <factory id="nxRuntimeFactory" class="org.nuxeo.ecm.platform.ui.granite.factory.NuxeoRuntimeServiceFact
     </factories >
```

Default channel

This is the default channel you have to use in your service-config.xml

```
    <channels >
         <channel-definition id="nx-amf" class="mx.messaging.channels.AMFChannel" >
             <endpoint
                uri="http://{server.name}:{server.port}/nuxeo/nuxeo-amf/amf"
                class="flex.messaging.endpoints.AMFEndpoint"/ >
         </channel-definition >
         <!--server.port and server.name are resolved dynamically at run time-- >
     </channels >
```

Services configuration is lighter on server side and same as usual on client side.

Seam Component

- client side

```
<services>
        <service
            id="flexDocumentManager"
            class="flex.messaging.services.RemotingService"
            messageTypes="flex.messaging.messages.RemotingMessage">
            <destination id="flexDocumentManager">
            <channels>
                <channel ref="nx-amf"/>
            </channels>
                <properties>
                    <factory>seamFactory</factory>
                    <source>flexDocumentManager</source>
                </properties>
            </destination>
        </service>
    </services>
```

- sever side

  Here is the minimum configuration to register your service using *NxGraniteConfigService* extension point:

```
<extension target="org.nuxeo.ecm.platform.ui.granite.config.NxGraniteConfigService" point="services">
       <seam id="myID" />
</extension>
```

  The id is the default service id. They will be used as the destination name if not specified. The id will be use a the default Seam component Name if not specified.

  A complete service declaration would look like this:

```
<extension target="org.nuxeo.ecm.platform.ui.granite.config.NxGraniteConfigService" point="services">
       <seam id="myID" destinationId="myDestinationId" source="mySeamComponentName"/>
</extension>
```

Runtime Services

- client side

```
<services>
        <service
            id="schemaManager"
            class="flex.messaging.services.RemotingService"
            messageTypes="flex.messaging.messages.RemotingMessage">
            <destination id="schemaManager">
            <channels>
                <channel ref="nx-amf"/>
            </channels>
                <properties>
                    <factory>nxruntimeFactory</factory>
                    <class>org.nuxeo.ecm.core.schema.SchemaManager</class>
                </properties>
            </destination>
        </service>
    </services>
```

- server side

  Here is the minimum configuration to register your service using *NxGraniteConfigService* extension point:

```
<extension target="org.nuxeo.ecm.platform.ui.granite.config.NxGraniteConfigService" point="services">
       <runtime id="myID" class="my.package.MyClass" />
</extension>
```

  The id is the default service id. They will be used as the destination name if not specified.

  A complete service declaration would look like this:

```
<extension target="org.nuxeo.ecm.platform.ui.granite.config.NxGraniteConfigService" point="services">
       <runtime id="myID" destinationId="myDestinationId" class="my.package.MyRuntimeServiceClass" />
</extension>
```

37.4.2. Your First Flex application

This the code of the facet explorer sample. The RemoteOject tag is binded to the schemaManager service by the destination attribute. For more information about the other tags, see Flex Documentation here : http://livedocs.adobe.com/flex/3/langref/index.html

```
<mx:Application xmlns:mx="http://www.adobe.com/2006/mxml"   xmlns:word="*">
        <mx:RemoteObject id="schemaManager" destination="schemaManager"/>
        <mx:Panel title="Nuxeo Facet explorer"
            paddingTop="10" paddingBottom="10" paddingLeft="10" paddingRight="10">
            <mx:VBox>
              <mx:TextInput id="facet" maxChars="30" text="Folderish"/>
              <mx:Button label="get Document Type"
                click="schemaManager.getDocumentTypeNamesForFacet(facet.text)"/>
              <mx:List id="typeList" height="100%" width="100%"
                dataProvider="{schemaManager.getDocumentTypeNamesForFacet.lastResult}" />
            </mx:VBox>
      </mx:Panel>
</mx:Application>
```

37.4.3. Granite DS configuration

- Externalizer

  If you want to use Nuxeo API, you will need a mapping between Java Object and Action Script Object. GraniteDS provides pluggable externalizer for your different Object. It aims to (de)serialize the different fields of ypur objects. For more information, see GraniteDS documentation here : http://www.graniteds.org/confluence/display/DOC/2.+Externalizers

- InvocationListener

  You might need more controil on mapping. For instance, the DocumentModel object in Nuxeo is rather complicated. So we have a FlexDocumentModel object wich is a simplified version of DocumentModel. The mapping between those two java objects is done in the NuxeoInvocationListener. It listens to each service invocation method call. Then we can switch from FlexDocumentModel to DocumentModel or the other way around. For more information on InvocationListener, see GraniteDS documentation here: http://www.graniteds.org/confluence/display/DOC/8.+Miscanellous+Options

Part V. Administration overview

In this part you will find some information and tips about common administration tasks on a Nuxeo EP server. If you need more details or specific cases, look at the administration guide.

Chapter 38. OS requirements, existing and recommended configuration

This chapter presents information about the running environment. Listing all required software, giving a recommended configuration and listing some others, known as operational, this chapter aims to help you to validate or define your production environment but the list is not exhaustive and needs to be completed with the users' experience.

38.1. Required software

- JBoss application server with EJB3 support enabled
- Java Development Kit (JDK)

38.2. Recommended configuration

38.2.1. Hardware configuration

Nuxeo EP is designed to be scalable and thus to be deployed on many servers. It can be installed on only one server for a start, and can also easily be installed on many servers. The constant is that there is the need to have one high-end server with good performances. Then the other servers can be more lower-end.

So the numbers below are given for the one needed high-end server.

- RAM: 2Gb is the minimum requirement for using Nuxeo EP
- CPU: Intel Core2 or equivalent and upper

  You might be better avoiding machines from the Intel Pentium 4 Xeon series since some models have a too small amount of cache. This impairs performance greatly compared to other CPU architecture of the same generation. Intel Pentium 4 servers are quite widespread because of an attractive price politic.
- Storage (disk) space: the minimum Nuxeo installation, along with the needed JBoss and libs, takes something between 200Mb and 250Mb on a filesystem. Then the final size will of course depend on the amount of data that will be stored in Nuxeo. A safe bet (until we provide better numbers) is to consider data space ratio of 1.5 to 2.

38.2.2. Default configuration

The default persistence configuration is lightweight and easy to use, but it is not made for performance.

- default Nuxeo 5.1 uses:
  - HSQL for SQL Data (directories, JBPM, Relations ...)
  - FileSystem persistence for Document repository
- default Nuxeo 5.2 uses:
  - Derby for SQL Data (directories, JBPM, Relations ...)
  - FileSystem persistence for Document repository

38.2.3. For optimal performances

- Linux
- PostgreSQL 8.2

  Use PostgreSQL for document repository and all other services except for Compass search engine (nxsearch-compass-ds.xml) which should be set to use filesystem.

  Configure the document repository to externalize the blobs to filesystem.

38.3. Known working configurations

38.3.1. OS

- Debian GNU/Linux 5.0 Lenny
- Linux Ubuntu 32 and 64 bits, Edgy, Feisty and Hardy (8.04).
- Linux Mandriva 2008.1
- Unix
- Mac OS X 10.4, 10.5
- Ms Windows 2003 server 32 and 64 bits, Windows XP

38.3.2. JVM

- Sun JDK 1.5 update 14, 15, 16

38.3.3. Storage backends

Different backends may be set as well for Nuxeo Core repository as for all other nuxeo services that persist data. See Chapter 40, *RDBMS Storage and Database Configuration* for more details, here is a list of known working backends.

- PostgreSQL 8.2
- PostgreSQL 8.3

  This version need a workaround to be applied as it is much stricter than PostgreSQL 8.2 with respect to value casting.

  Execute the following commands in your PostgreSQL console:

```
CREATE FUNCTION pg_catalog.text(integer) RETURNS text STRICT IMMUTABLE LANGUAGE SQL AS 'SELECT textin(int4out($1
CREATE CAST (integer AS text) WITH FUNCTION pg_catalog.text(integer) AS IMPLICIT;
COMMENT ON FUNCTION pg_catalog.text(integer) IS 'convert integer to text';

CREATE FUNCTION pg_catalog.text(bigint) RETURNS text STRICT IMMUTABLE LANGUAGE SQL AS 'SELECT textin(int8out($1)
CREATE CAST (bigint AS text) WITH FUNCTION pg_catalog.text(bigint) AS IMPLICIT;
COMMENT ON FUNCTION pg_catalog.text(bigint) IS 'convert bigint to text';
```

  See FAQ about using PostgreSQL 8.3.

- MySQL
- Oracle 10
- MsSQL 2005

- HSQL

  Default Nuxeo 5.1 embedded database.

- Derby

  Default Nuxeo 5.2 embedded database.

38.3.4. LDAP

- OpenLDAP
- Ms Active directory

Chapter 39. SMTP Server configuration

On the Nuxeo EP built-in types, you can manage e-mailing and notifications. Before getting this features working, you need to configure your SMTP server. Nuxeo EP relies on the application server mail service for mailing stuff. So you just need to configure the `mail-service.xml` file in `$JBOSS_HOME/server/default/deploy/`. You can find examples of how to use this file in the JBoss wiki, and detailed information about the properties of this file in the JavaMail Javadoc.

Chapter 40. RDBMS Storage and Database Configuration

To run Nuxeo EP for real-world application, you need to get rid of the default embedded database and set up a real RDBMS server (such as PostgreSQL, MySQL, Oracle, etc.) to store Nuxeo EP's data.

In order to define a SQL DB as repository back-end you have to :

- install the JDBC driver for your DBMS,
- configure the Nuxeo Core storage, modifying the default repository configuration,
- configure your database.
- Eventually, configure storage for other Nuxeo services.
- You may also add a new repository configuration (and optionally disable the old one).

40.1. Storages in Nuxeo EP

Nuxeo EP manages several types of data: Documents, relations, audit trail, users, groups ...

Each type of data has its own storage engine and can be configured separately. All storages can use RDBMS but some of them can also use the filesystem.

This means you can have a RDBMS only configuration or a mixed configuration using RDBMS and filesystem. You can even use several RDBMS if you find a use case for that.

For a lot of services, RDBMS access are encapsulated by JPA or hibernate calls, so the storage should be RDBMS agnostic as long as there is a hibernate dialect for the DB.

40.2. Installing the JDBC driver

To enable the connection to the database, you first need to install the JDBC driver into `$JBOSS_HOME/server/default/lib/`.

Here are some drivers download locations:

- PostgreSQL JDBC Drivers
- MySQL JDBC Drivers

40.3. Configuring Nuxeo Core Storage

Nuxeo Core stores data for the documents themselves: the hierarchy of documents, their metadata and security, and the attached files. There are two main Nuxeo Core Storage backends: the recent (available in Nuxeo 5.2) Visible Contents Store (VCS), based on a mapper to native RDBMS tables, and the previous JCR-based backend, using Jackrabbit.

This section will show you how to configure each backend, using PostgreSQL 8.3 as an example underlying storage. The setup for other RDBMS should be very similar.

40.3.1. Visible Content Store configuration

To set up VCS, you first need create a datasource for it. The datasource is *not* a standard JDBC datasource, so

has different syntax, even though it contains information about JDBC connection parameters. This file is usually named
`$JBOSS_HOME/server/default/deploy/nuxeo.ear/datasources/default-repository-ds.xml`.

Example 40.1. Datasource for VCS using PostgreSQL

```
<?xml version="1.0"?>
<connection-factories>
  <tx-connection-factory>
    <jndi-name>NXRepository/default</jndi-name>
    <xa-transaction/>
    <track-connection-by-tx/>
    <adapter-display-name>Nuxeo SQL Repository DataSource</adapter-display-name>
    <rar-name>nuxeo.ear#nuxeo-core-storage-sql-ra-1.5-SNAPSHOT.rar</rar-name>
    <connection-definition>org.nuxeo.ecm.core.storage.sql.Repository</connection-definition>
    <config-property name="name">default</config-property>
    <config-property name="xaDataSource" type="java.lang.String">org.postgresql.xa.PGXADataSource</config-proper
    <config-property name="property" type="java.lang.String">ServerName=localhost</config-property>
    <config-property name="property" type="java.lang.String">PortNumber/Integer=5432</config-property>
    <config-property name="property" type="java.lang.String">DatabaseName=nuxeo</config-property>
    <config-property name="property" type="java.lang.String">User=nuxeo</config-property>
    <config-property name="property" type="java.lang.String">Password=password</config-property>
    <max-pool-size>20</max-pool-size>
  </tx-connection-factory>
</connection-factories>
```

You will then need to specify the actual repository configuration, usually store in a file named
`$JBOSS_HOME/server/default/deploy/nuxeo.ear/config/default-repository-config.xml`

Example 40.2. Repository Configuration for VCS

```
<?xml version="1.0"?>
<component name="default-repository-config">
  <extension target="org.nuxeo.ecm.core.repository.RepositoryService" point="repository">
    <repository name="default"
      factory="org.nuxeo.ecm.core.storage.sql.coremodel.SQLRepositoryFactory">
      <repository name="default">
        <indexing>
          <!-- example configuration for H2
          <fulltext analyzer="org.apache.lucene.analysis.fr.FrenchAnalyzer"/>
          -->
          <!-- example configuration for PostgreSQL
          <fulltext analyzer="french"/>
          -->
          <!-- example configuration for Microsoft SQL Server
          <fulltext catalog="nuxeo" analyzer="french"/>
          -->
        </indexing>
        <!-- uncomment this to enable clustering
             delay is in milliseconds
               default delay is 0 (no delay before processing invalidations)
        <clustering enabled="true" delay="1000" />
        -->
      </repository>
    </repository>
  </extension>
</component>
```

40.3.2. JCR backend configuration

First you have to specify a datasource in
`$JBOSS_HOME/server/default/deploy/nuxeo.ear/datasources/default-repository-ds.xml`.

Example 40.3. Datasource for JCR backend

```
<?xml version="1.0"?>
<!DOCTYPE connection-factories PUBLIC
  "-//JBoss//DTD JBOSS JCA Config 1.5//EN"
```

```
  "http://www.jboss.org/j2ee/dtd/jboss-ds_1_5.dtd">
<connection-factories>
  <mbean code="org.nuxeo.ecm.core.repository.JBossRepository"
         name="nx:type=repository,name=default">
    <constructor>
      <arg type="java.lang.String" value="default"/>
    </constructor>
  </mbean>
  <tx-connection-factory>
    <jndi-name>NXRepository/default</jndi-name>
    <adapter-display-name>NX Repository Adapter</adapter-display-name>
    <rar-name>nuxeo.ear#nuxeo-core-jca-${project.version}.rar</rar-name>
    <connection-definition>org.nuxeo.ecm.core.model.Repository</connection-definition>
    <xa-transaction/>
    <!-- Configuration properties. -->
    <config-property name="name">default</config-property>
  </tx-connection-factory>
</connection-factories>
```

Then edit `$JBOSS_HOME/server/default/deploy/nuxeo.ear/config/default-repository-config.xml`.

Example 40.4. Default repository configuration

```
<?xml version="1.0"?>
<component name="default-repository-config">
  <documentation>
    Defines the default JackRabbit repository used for development and testing.
  </documentation>
  <extension target="org.nuxeo.ecm.core.repository.RepositoryService"
    point="repository">
    <documentation>
      Declare a JackRabbit repository to be used for development and tests. The
      extension content is the Jackrabbit XML configuration of the repository.
    </documentation>
    <repository name="default"
      factory="org.nuxeo.ecm.core.repository.jcr.JCRRepositoryFactory"
      securityManager="org.nuxeo.ecm.core.repository.jcr.JCRSecurityManager"
      forceReloadTypes="false">
      <Repository>
        <!--
          virtual file system where the repository stores global state
          (e.g. registered namespaces, custom node types, etc.)
        -->
        <FileSystem class="org.apache.jackrabbit.core.fs.local.LocalFileSystem">
          <param name="path" value="${rep.home}/repository" />
        </FileSystem>
        <!--
          security configuration
        -->
        <Security appName="Jackrabbit">
          <!--
            access manager:
            class: FQN of class implementing the AccessManager interface
          -->
          <AccessManager class="org.apache.jackrabbit.core.security.SimpleAccessManager">
            <!-- <param name="config" value="${rep.home}/access.xml"/> -->
          </AccessManager>
          <LoginModule class="org.apache.jackrabbit.core.security.SimpleLoginModule">
            <!-- anonymous user name ('anonymous' is the default value) -->
            <param name="anonymousId" value="anonymous" />
            <!--
              default user name to be used instead of the anonymous user
              when no login credentials are provided (unset by default)
            -->
            <!-- <param name="defaultUserId" value="superuser"/> -->
          </LoginModule>
        </Security>

        <!--
          location of workspaces root directory and name of default workspace
        -->
        <Workspaces rootPath="${rep.home}/workspaces" defaultWorkspace="default" />
        <!--
          workspace configuration template:
          used to create the initial workspace if there's no workspace yet
        -->
        <Workspace name="${wsp.name}">
          <!--
            virtual file system of the workspace:
            class: FQN of class implementing the FileSystem interface
          -->
          <FileSystem class="org.apache.jackrabbit.core.fs.local.LocalFileSystem">
            <param name="path" value="${wsp.home}" />
          </FileSystem>

          <!--
```

```
            persistence manager of the workspace:
            class: FQN of class implementing the PersistenceManager interface
          -->
          <PersistenceManager class="org.apache.jackrabbit.core.state.obj.ObjectPersistenceManager">
          </PersistenceManager>

          <!--
            Search index and the file system it uses.
            class: FQN of class implementing the QueryHandler interface
          -->
          <SearchIndex class="org.apache.jackrabbit.core.query.lucene.SearchIndex">
            <param name="path" value="${wsp.home}/index" />
          </SearchIndex>
        </Workspace>

        <!--
          Configures the versioning
        -->
        <Versioning rootPath="${rep.home}/version">
          <!--
            Configures the filesystem to use for versioning for the respective
            persistence manager
          -->
          <FileSystem class="org.apache.jackrabbit.core.fs.local.LocalFileSystem">
            <param name="path" value="${rep.home}/version" />
          </FileSystem>

          <!--
            Configures the persistence manager to be used for persisting version state.
            Please note that the current versioning implementation is based on
            a 'normal' persistence manager, but this could change in future
            implementations.
          -->
          <PersistenceManager class="org.apache.jackrabbit.core.state.obj.ObjectPersistenceManager">
          </PersistenceManager>
        </Versioning>

        <!--
          Search index for content that is shared repository wide
          (/jcr:system tree, contains mainly versions)
        -->
        <SearchIndex class="org.apache.jackrabbit.core.query.lucene.SearchIndex">
          <param name="path" value="${rep.home}/repository/index" />
        </SearchIndex>
      </Repository>
    </repository>
  </extension>
</component>
```

Change the two `PersistenceManager` sections defining various database connection settings.

Refer to the Jackrabbit documentation for more information, and to the Jackrabbit Javadoc for details about configuring the PersistenceManager.

In particular, decide if you want the binary blobs stored inside the database or in the filesystem (change `externalBLOBs` to `true` if you want them outside the database, in the filesystem).

Using externalized Blobs can provide a performance improvement in particular if you need to store a lot of big files. Depending on the RDBMS used, there may also be a max size limit for blob if you store them in the RDBMS (for PostgeSQL 8.2 blobs are limited to 1 GB).

Here are some examples:

Example 40.5. Sample configuration for a PostgreSQL Jackrabbit repository

```
        <!-- Workspaces configuration. Nuxeo only uses the default workspace. -->
        (...)
          <PersistenceManager class="org.apache.jackrabbit.core.persistence.bundle.PostgreSQLPersistenceManager"
            <param name="driver" value="org.postgresql.Driver"/>
            <param name="url" value="jdbc:postgresql://localhost/nuxeo"/>
            <param name="user" value="postgres"/>
            <param name="password" value="password"/>
            <param name="schema" value="postgresql"/>
            <param name="schemaObjectPrefix" value="jcr_${wsp.name}_"/>
            <param name="externalBLOBs" value="false"/>
          </PersistenceManager>
        (...)
        <!-- Versioning configuration. -->
```

```
  (...)
  <PersistenceManager class="org.apache.jackrabbit.core.persistence.bundle.PostgreSQLPersistenceManager"
    <param name="driver" value="org.postgresql.Driver"/>
    <param name="url" value="jdbc:postgresql://localhost/nuxeo"/>
    <param name="user" value="postgres"/>
    <param name="password" value="password"/>
    <param name="schema" value="postgresql"/>
    <param name="schemaObjectPrefix" value="jcr_ver_"/>
    <param name="externalBLOBs" value="false"/>
  </PersistenceManager>
```

Example 40.6. Sample configuration for a MySQL Jackrabbit repository

```
<!-- Workspaces configuration. Nuxeo only uses the default workspace. -->
(...)
  <PersistenceManager class="org.apache.jackrabbit.core.persistence.bundle.MySqlPersistenceManager">
    <param name="driver" value="com.mysql.jdbc.Driver"/>
    <param name="url" value="jdbc:mysql://localhost/nuxeo"/>
    <param name="user" value="mysql"/>
    <param name="password" value="password"/>
    <param name="schema" value="mysql"/>
    <param name="schemaObjectPrefix" value="jcr_${wsp.name}_"/>
    <param name="externalBLOBs" value="true"/>
  </PersistenceManager>
(...)
<!-- Versioning configuration. -->
  (...)
  <PersistenceManager class="org.apache.jackrabbit.core.persistence.bundle.MySqlPersistenceManager">
    <param name="driver" value="com.mysql.jdbc.Driver"/>
    <param name="url" value="jdbc:mysql://localhost/nuxeo"/>
    <param name="user" value="mysql"/>
    <param name="password" value="password"/>
    <param name="schema" value="mysql"/>
    <param name="schemaObjectPrefix" value="jcr_ver_"/>
    <param name="externalBLOBs" value="true"/>
  </PersistenceManager>
```

Note: "schemaObjectPrefix" must have different values in workspace & versioning configuration

For JackRabbit, there are some persistence manager specific to each RDBMS :

- for PostgreSQL: you may use org.apache.jackrabbit.core.persistence.bundle.PostgreSQLPersistenceManager
- for MySQL: you may use org.apache.jackrabbit.core.persistence.bundle.MySqlPersistenceManager
- for Oracle 10: you may use org.apache.jackrabbit.core.persistence.bundle.OraclePersistenceManager
- for MSSQL2005 you may use : org.apache.jackrabbit.core.persistence.bundle.MSSqlPersistenceManager

40.3.3. Set up your RDBMS

Create the database in the database server, enable IP connection, setup permissions and test the connection.

40.3.4. Start Nuxeo EP

You can now start JBoss AS and verify that your new repository is used!

40.4. Configuring Storage for other Nuxeo Services

Many services beyond Nuxeo Core Repository are using an SQL database to persist their data, such as:

- Relations Service: RDF data is stored in SQL,
- Audit Service: Audit logs are stored via JPA,
- Directories: entries can be stored into an SQL database.

By default, all these services use the JBoss AS's embedded HSQLDB.

40.4.1. Configuring datasources

Each service can use a dedicated datasource to define the database connection. However, to simplify configuration, all datasources are JNDI NamingAlias pointing to a single datasource (NuxeoDS). If you want to change the default database, you can simply do the following:

- deploy the needed JDBC Driver in `$JBOSS_HOME/server/default/lib`,
- modify the datasource definition file in `$JBOSS_HOME/server/default/deploy/nuxeo.ear/datasources/unified-nuxeo-ds.xml`.

The default configuration has a commented out PostgresSQL configuration.

To use a dedicated datasource for a service, you need to modify its configuration file. If you would like to store audit logs in PostgreSQL using its own datasource, remove the NamingAlias in `nxaudit-logs-ds.xml` and replace it with the datasource configuration example:

Example 40.7. Datasource for the Audit Service using PostgreSQL

```
<?xml version="1.0"?>
<datasources>
  <local-tx-datasource>
    <jndi-name>nxaudit-logs</jndi-name>
    <connection-url>jdbc:postgresql://localhost/logs</connection-url>
    <driver-class>org.postgresql.Driver</driver-class>
    <user-name>username</user-name>
    <password>password</password>
  </local-tx-datasource>
</datasources>
```

We recommend to enable XA transactions if your database server support it (for PostgreSQL, you have to use 8.x versions). The following datasource definition example enables XA transaction for the Audit Service using PostgreSQL.

Example 40.8. Datasource for the Audit Service using PostgreSQL with XA transactions

```
<?xml version="1.0"?>
<datasources>
   <xa-datasource>
     <jndi-name>nxaudit-logs</jndi-name>
     <xa-datasource-class>org.postgresql.xa.PGXADataSource</xa-datasource-class>
     <xa-datasource-property name="ServerName">localhost</xa-datasource-property>
     <xa-datasource-property name="PortNumber">5432</xa-datasource-property>
     <xa-datasource-property name="DatabaseName">logs</xa-datasource-property>
     <xa-datasource-property name="User">postgres</xa-datasource-property>
     <xa-datasource-property name="Password">password</xa-datasource-property>
     <track-connection-by-tx/>
   </xa-datasource>
</datasources>
```

See Datasources Configuration on the JBoss wiki for more examples of datasources.

This method works for most services:

- Audit: `nxaudit-logs-ds.xml`
- Placeful Configuration Service & Subscriptions: `nxplaceful-ds.xml`
- UID generator: `nxuidsequencer-ds.xml`
- jBPM engine: `nxworkflow-jbpm-ds.xml`
- Workflow Document Service: `nxworkflow-documents-ds.xml`
- Archive management: `nxarchive-records-ds.xml`
- Relations: `nxrelations-default-jena-ds.xml`
- Compass search engine: `nxsearch-compass-ds.xml`

40.4.1.1. Special MySQL needs

MySQL is a quirky database which sometimes needs very specific options to function properly.

The datasources used by Jena (`nxrelations-default-jena-ds.xml` and `nxcomment-jena-ds.xml`) need to use a "relax autocommit" mode. To enable that, change the `connection-url` in the datasources to something like:

```
<connection-url>
  jdbc:mysql://localhost/nuxeo?relaxAutoCommit=true
</connection-url>
```

The datasource used by Compass (`nxsearch-compass-ds.xml`) needs to be put in "relax autocommit" too, and in addition it needs an "emulate locators" option:

```
<connection-url>
  jdbc:mysql://localhost/nuxeo?relaxAutoCommit=true&amp;emulateLocators=true
</connection-url>
```

This is documented at http://static.compassframework.org/docs/latest/jdbcdirectory.html .

Note the `&` syntax for the URL parameters instead of just `&` because the URL is embedded in an XML file.

40.4.2. Relation service configuration

The Relation Service uses a datasource to define the data storage. However modifying the datasource is not enough, you also have to tell to the Jena engine which database dialect is used, as it doesn't auto-detect it.

To do that, edit the `sql.properties` file in `$JBOSS_HOME/server/default/deploy/nuxeo.ear/config/` and change the definition of the `org.nuxeo.ecm.sql.jena.databaseType` property. The possible values are:

- `Derby`
- `HSQL`
- `MsSQL`
- `MySQL`
- `Oracle`
- `PostgreSQL`

In the same file, the property `org.nuxeo.ecm.sql.jena.databaseTransactionEnabled` is `false` by default, but must be set to `true` for Oracle.

Please refer to the Jena Site for more information about database support.

The value of the above properties are used as variables by the extension point defining the Jena configuration, so that they only have to be changed in one place.

40.4.3. Compass search engine dialect configuration

Note that this section is obsolete in Nuxeo 5.2, and should not be used anymore.

The Compass plugin is configured using a datasource, but at the time of this writing it still needs some additional configuration in a file embedded in its Jar. You should go to `$JBOSS_HOME/server/default/deploy/nuxeo.ear/system/` and inside the directory `nuxeo-platform-search-compass-plugin-5.1-SNAPSHOT.jar` (the version number may be different) then edit the file `compass.cfg.xml`. You will find a section like:

```
<connection>
  <jdbc managed="true"
    dialectClass="org.apache.lucene.store.jdbc.dialect.HSQLDialect"
    deleteMarkDeletedDelta="3600000">
    <dataSourceProvider>
      <jndi lookup="java:/nxsearch-compass" />
    </dataSourceProvider>
  </jdbc>
</connection>
```

The `dialectClass` has to be changed according to your datasource configuration. The possible values end with `MySQLDialect`, `PostgreSQLDialect` , etc. They are documented in the Compass documention about SQL dialects .

40.5. Setting up a new repository configuration

If you just want to change the `default` repository name appearing in the url `http://.../nuxeo/nxdoc/default/`, modify the repository name value in:

- `default-repository-config.xml`
- `platform-config.xml`

TODO: Nuxeo configuration has changed, the two above sections need to be updated.

40.5.1. Add the new repository configuration

Create a repository definition as a contribution to the extension point `org.nuxeo.ecm.core.repository.RepositoryService` (for example: `MyRepo-repositoy-config.xml`) in `$JBOSS_HOME/server/default/deploy/nuxeo.ear/config/`.

You can take example on the default repository definition $JBOSS_HOME/server/default/deploy/nuxeo.ear/config/default-repository-config.xml.

You have to properly configure the following aspects:

- the name of the component (`name="org.nuxeo.project.sample.repository.MyRepo"`), which must be unique among component names,
- the name of the repository (`<repository name="MyRepo">`), which is used to refer to it from your

application and must also be unique among repository names,

- the various database connection settings (driver, user, password, schema, etc.),
- decide if you want the binary blobs stored inside the database or in the filesystem (change `externalBLOBs` to `true` if you want them outside the database).

Refer to the Jackrabbit documentation for more information, and to the Jackrabbit Javadoc for details about configuring the `BundleDbPersistenceManager`.

Example 40.9. Sample configuration for a PostgreSQL Jackrabbit repository

```
<?xml version="1.0"?>
<component name="org.nuxeo.project.sample.repository.MyRepo">
  <extension target="org.nuxeo.ecm.core.repository.RepositoryService" point="repository">
    <repository name="MyRepo"
                factory="org.nuxeo.ecm.core.repository.jcr.JCRRepositoryFactory"
                securityManager="org.nuxeo.ecm.core.repository.jcr.JCRSecurityManager"
                forceReloadTypes="false">
      <Repository>
        <FileSystem class="org.apache.jackrabbit.core.fs.local.LocalFileSystem">
          <param name="path" value="${rep.home}/repository"/>
        </FileSystem>
        <Security appName="Jackrabbit">
          <AccessManager class="org.apache.jackrabbit.core.security.SimpleAccessManager">
        </AccessManager>
          <LoginModule class="org.apache.jackrabbit.core.security.SimpleLoginModule">
            <param name="anonymousId" value="anonymous"/>
          </LoginModule>
        </Security>

        <!-- Workspaces configuration. Nuxeo only uses the default workspace. -->
        <Workspaces rootPath="${rep.home}/workspaces" defaultWorkspace="default"/>
        <Workspace name="${wsp.name}">
          <FileSystem class="org.apache.jackrabbit.core.fs.local.LocalFileSystem">
            <param name="path" value="${wsp.home}"/>
          </FileSystem>
          <PersistenceManager class="org.apache.jackrabbit.core.persistence.bundle.PostgreSQLPersistenceManager"
            <param name="driver" value="org.postgresql.Driver"/>
            <param name="url" value="jdbc:postgresql://localhost/nuxeo"/>
            <param name="user" value="postgres"/>
            <param name="password" value="password"/>
            <param name="schema" value="postgresql"/>
            <param name="schemaObjectPrefix" value="jcr_${wsp.name}_"/>
            <param name="externalBLOBs" value="false"/>
          </PersistenceManager>
          <SearchIndex class="org.apache.jackrabbit.core.query.lucene.SearchIndex">
            <param name="path" value="${wsp.home}/index"/>
          </SearchIndex>
        </Workspace>

        <!-- Versioning configuration. -->
        <Versioning rootPath="${rep.home}/version">
          <FileSystem class="org.apache.jackrabbit.core.fs.local.LocalFileSystem">
            <param name="path" value="${rep.home}/version"/>
          </FileSystem>
          <PersistenceManager class="org.apache.jackrabbit.core.persistence.bundle.PostgreSQLPersistenceManager"
            <param name="driver" value="org.postgresql.Driver"/>
            <param name="url" value="jdbc:postgresql://localhost/nuxeo"/>
            <param name="user" value="postgres"/>
            <param name="password" value="password"/>
            <param name="schema" value="postgresql"/>
            <param name="schemaObjectPrefix" value="jcr_ver_"/>
            <param name="externalBLOBs" value="false"/>
          </PersistenceManager>
        </Versioning>

        <!-- Index for repository-wide information, mainly versions. -->
        <SearchIndex class="org.apache.jackrabbit.core.query.lucene.SearchIndex">
          <param name="path" value="${rep.home}/repository/index"/>
        </SearchIndex>
      </Repository>
    </repository>
  </extension>
</component>
```

40.5.2. Declare the new repository to the platform

TODO: this should be moved to a different section as it doesn't pertain to the SQL configuration itself.

You have now replaced the default repository (`demo`) with your newly defined one (`MyRepo`). To actually use it, create or edit the file `MyPlatform-Layout-config.xml` in `$JBOSS_HOME/server/default/deploy/nuxeo.ear/config/` and configure the parameters as shown in the following example.

```
<?xml version="1.0"?>
<component name="MyPlatformLayout">

  <require>org.nuxeo.ecm.platform.api.DefaultPlatform</require>

  <extension target="org.nuxeo.ecm.platform.util.LocationManagerService" point="location">
    <locationManagerPlugin> <!-- This disable the default (demo) repository -->
      <locationEnabled>false</locationEnabled>
      <locationName>demo</locationName>
    </locationManagerPlugin>
    <locationManagerPlugin> <!-- This enable your new repository -->
      <locationEnabled>true</locationEnabled>
      <locationName>MyRepo</locationName> <!-- Use the name of your repository -->
    </locationManagerPlugin>
  </extension>

 <extension target="org.nuxeo.ecm.core.api.repository.RepositoryManager"
    point="repositories">
    <documentation>The default repository</documentation>
    <repository name="MyRepo" label="My Repository"/>
  </extension>

  <extension target="org.nuxeo.runtime.api.ServiceManagement" point="services">
      <service class="org.nuxeo.ecm.core.api.CoreSession" name="MyRepo" group="core">
          <locator>%DocumentManagerBean</locator>
      </service>
  </extension>

</component>
```

This sample configuration creates a new repository in the `core` Group that will be assigned to the default server. If you want to have it located on an other server you can use:

```

<extension target="org.nuxeo.runtime.api.ServiceManagement" point="servers">
  <!-- define new locator for group MyGroup -->
  <server class="org.nuxeo.runtime.api.JBossServiceLocator">
    <group>MyGroup</group>
    <property name="java.naming.factory.initial">org.jnp.interfaces.NamingContextFactory</property>
    <property name="java.naming.provider.url">jnp://MyServer:1099</property>
    <property name="java.naming.factory.url.pkgs">org.jboss.naming:org.jnp.interfaces</property>
  </server>

  <!-- bind MyRepo to MyGroup -->
  <extension target="org.nuxeo.runtime.api.ServiceManagement" point="services">
    <service class="org.nuxeo.ecm.core.api.CoreSession" name="MyRepo" group="MyGroup">
      <locator>%DocumentManagerBean</locator>
    </service>
  </extension>

</extension>
```

Chapter 41. LDAP Integration

41.1. For users/groups storage backend

The user interface in Nuxeo EP gets the data from NXDirectory. As a consequence you can choose your source. By default, the users/groups data is stored in a SQL database. If you want to get the users from a LDAP directory, you need to deploy one of the following configuration:

- Users in LDAP, groups in SQL

  Go to the `examples` sub-folder and copy the `default-ldap-users-directory-bundle.xml` file in the `nuxeo.ear/config` folder of the JBoss instance (or bundle it in a jar, cf packaging in this guide). This sample setup replaces the default userDirectory configuration SQL with users fetched from the LDAP server. The groupDirectory remains unaffected by this setup. You might want to copy the file `default-virtual-groups-bundle.xml` and adjust defaultAdministratorId to select a user from your LDAP that have administrative rights by default. You can also configure the section on defaultGroup to make all users members of some default group (typically the members group) so that they have default right without having to make them belong to groups explicitly.

- Users and groups in LDAP

  Copy the users setup as previously; moreover copy the `default-ldap-groups-directory-bundle.xml` file in the `nuxeo.ear/config` folder of the JBoss instance. This sample setup which is dependent on the previous one additionally overrides the default `groupDirectory` setup to read the groups from the LDAP directory typically from `groupOfUniqueNames` entries with fully qualified `dn` references to the user entries or to subgroups. You can edit the `nuxeo.ear/config/*.xml` files on the JBoss instance, but you will need to restart JBoss to take changes into account.

Chapter 42. OpenOffice.org server installation

OpenOffice.org (OOo) is used on server-side for different tasks such as file transforms (eg. to PDF) or other advanced tasks.

42.1. Installation

See http://www.openoffice.org/ for installation procedure if not provided by your OS.

42.1.1. Start server

Since version 2.3, OpenOffice can be started in headless mode. This means that under linux, you non longer need to run a Xvfb.

Use the oooctl control script or, depending on your system and installation method, start OOo running :

- for Linux:

```
/path/to/openoffice/program/soffice.bin
            -headless -nofirststartwizard
            -accept="socket,host=localhost,port=8100;urp;StarOffice.Service"
```

- for Mac OS X:

```
/path/to/openoffice.app/Contents/MacOS/soffice.bin
            -headless -nofirststartwizard
            -accept="socket,host=localhost,port=8100;urp;StarOffice.Service"
```

- for Windows:

```
soffice.exe -headless
            -nofirststartwizard
            -accept="socket,host=localhost,port=8100;urp;StarOffice.Service"
```

42.1.2. Parameters

42.1.2.1. `-headless`

The `-headless` switch lets OOo manage to answer every question that may occur with a default response avoiding dead-locks. It also hides the OOo windows if a display is available by default (eg. Ms Windows). Note that the `-invisible` switch is not used as it is redundant and errors on PDF exports may occur.

42.1.2.2. `-nofirststartwizard`

The `-nofirststartwizard` skips any post-install wizard, allowing to use OOo directly after the installation without any user parameterization.

Install nxSkipInstallWizard.oxt extension to make this parameter permanent.

42.1.2.3. `-accept`

The `UNO` protocol implies that OOo is opened in listening mode on a network interface. This parameter let OOo listen on the right interface and port.

Install nxOOoAutoListen.oxt extension to set this listening permanent.

Alternate method: this could also be done by adding this connection info in OOo `Setup.xcu` configuration file.

```
<node oor:name="Office">
  <prop oor:name="ooSetupConnectionURL">
    <value>socket,host=localhost,port=8100;urp;StarOffice.Service</value>
  </prop>
</node>
```

42.1.3. Installing an extension

Nuxeo has developed some tools to ease the settings, they are available at the Nuxeo svn tools ooo section as extensions to install in OpenOffice. This can be done thru the GUI via the menu "Tools>Extensions Manager>Add..." or by opening the wanted extension with OpenOffice. To do it from a command line, run:

```
/path/to/openoffice/program/unopkg add extension.oxt
```

42.1.4. Notes

42.1.4.1. Multi-machine

On multi-machine deployment, we recommend to install OOo on the server running the webapp (ie stateless server on a bi-machine "stateless/stateful" installation).

42.1.4.2. OpenOffice lower than 2.3.x

For OOo versions lower than 2.3, a graphical interface is needed. If the server that hosts OOo is Linux server that has no X installed, then the X virtual frame buffer tool `Xvfb` (or `xvfb-run` depending of your distribution) can be used to create a suitable display.

```
Xvfb :77 -auth Xperm -screen 0 1024x768x24
export DISPLAY=":77.0"
```

```
xvfb-run -a /path/to/openoffice/program/soffice.bin
        -headless -nofirststartwizard
        -accept="socket,host=localhost,port=8100;urp;StarOffice.Service"
```

42.1.4.3. 32/64 bits

Note that the platform used for both the JVM and OOo have to be consistent to allow UNO protocol to work: with a 32 bits JVM, you'll have to use the 32 bits OOo version while the 64 bits one will be mandatory for a 64 bits JVM.

42.2. Running OpenOffice as a Daemon

42.2.1. Nuxeo OOo Daemon

Starting with 5.2-M4, Nuxeo includes a Daemon to start and manage OpenOffice server.

This daemon is based on the OpenOffice Server Daemon project

OOo Daemon provides :

- OpenOffice server automatic startup

Depending on configuration the OOo Server may be started automatically when Nuxeo starts or on first call.

- OpenOffice server instances pooling

  Daemon can manage several OpenOffice instance (workers) and distribute tasks across them. In order to avoid leaks, OOo workers are recycled.

42.2.2. Configuring Nuxeo OOo Daemon

Configuration is done via an xml extension point. Defaut config can be editer in file located in `nuxeo.ear/config/ooo-config.xml`.

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.platform.convert.oooDaemon.config.default">

  <extension target="org.nuxeo.ecm.platform.convert.oooserver.OOoDaemonManagerComponent"
    point="oooServerConfig">

  <OOoServer>
    <!-- enable Nuxeo Daemon to manage OOo server instances : default is true -->
    <enableDaemon>true</enableDaemon>
    <!-- define OOo server listen IP : used even if daemon is disabled -->
    <oooListenIP>127.0.0.1</oooListenIP>
    <!-- define OOo server listen port : used even if daemon is disabled -->
    <oooListenPort>8100</oooListenPort>
    <!-- define Daemon listen port : used only if daemon is enabled -->
    <oooDaemonListenPort>8101</oooDaemonListenPort>
    <!-- define number of OOo worker process : used only if daemon is enabled -->
    <oooWorkers>1</oooWorkers>

    <!-- define OOo installation path : used only if daemon is enabled -->
    <!-- if not defined Nuxeo will try to find the path automatically -->
    <!--<oooInstallationPath>/usr/lib/openoffice/program</oooInstallationPath>-->

    <!-- define jpipe library path : used only for OOo 3 -->
    <!--<jpipeLibPath>/opt/openoffice.org/ure/lib</jpipeLibPath>-->

    <!-- define number of time a worker process can be used before being recycled: used only if daemon is enable
    <oooWorkersRecycleInterval>10</oooWorkersRecycleInterval>

    <!-- define is Daemon is started at server startup : used only if daemon is enabled -->
    <autoStart>false</autoStart>

  </OOoServer>

  </extension>

</component>
```

In most cases, you should not have to define the location of your OpenOffice installation as the Daemon will try to find it in the standard installation path.

If you want to use OpenOffice 3.x, you have to specify the jpipeLibPath in order to enable jpipe that is used by the Daemon to communicate with OpenOffice workers.

Chapter 43. Run Nuxeo EP with a specific IP binding

To be able to call the Nuxeo Services remotely using the Nuxeo Framework (e.g. when using Nuxeo RCP), you will need to bind an IP address when running the server. To do this, a few step are needed:

- in `$JBOSS_HOME/server/default/deploy/nuxeo.ear/config/nuxeo.properties` change the value of `org.nuxeo.ecm.instance.host`, replace "localhost" by the IP address of the JBoss server,
- in `$JBOSS_HOME/server/default/deploy/nuxeo.ear/config/`platform-config.xml

  mono-server: ${org.nuxeo.ecm.instance.host} is used to reference address set in nuxeo.properties

  multi-server: either use IP or set DNS aliases in your hosts file (default names in Nuxeo packages are usually something like nxcoreserver, nxsearchserver, nxplatformserver, nxjmsserver, nxwebserver, nxdbserver).
- in `$JBOSS_HOME/server/default/deploy/nuxeo.ear/config/datasources/core-events-ds.xml` replace, if enabled, the 127.0.0.1 value of `java.naming.provider.url` by the IP address of the Jboss server,
- start jboss with the -b option:

  ./run.sh -b server_IP_address

IP behavior change between JBoss 4.0.x and JBoss 4.2.x

When no IP is specified, JBoss 4.0 is bound to listening on any address, not only on localhost.

For obvious security reasons, JBoss 4.2 (version required by Nuxeo EP 5.2) has a different behavior. If you still want this, use **-b 0.0.0.0**

For a server in production, see SecureJBoss.

Chapter 44. Backup, restore and reset

Note

For server migration (moving a Nuxeo instance from a server to another), follow the backup procedure on the source server and then the restore procedure on the destination server.

Note

Nuxeo-shell may be used for import/export purpose but it's not the same as the "system backup" described here.

44.1. Backup

Important

It is highly recommended to stop nuxeo during backup.

If you use a database, you may switch it in a "hot backup" mode to ensure that all new incoming requests will wait the end of the backup process.

Backup these directories and files:

- Nuxeo libraries deployed in JBoss as they are specific to the running Nuxeo EP version

  ```
  $JBOSS/server/default/lib/nuxeo*
  ```

- Nuxeo EP EAR (Entreprise ARchive)

  ```
  $JBOSS/server/default/deploy/nuxeo.ear/
  ```

- data stored on filesystem

  ```
  $JBOSS/server/default/data/
  ```

- it could be useful to backup logs

  ```
  $JBOSS/server/default/log/
  ```

Backup your database(s) if you use one (some).

Finally, backup any file you have customized. Here is a short list of such files :

- if you changed JVM startup parameters (run.bat, run.conf) or use Nuxeo's startup script (jbossctl, jbossctl.conf, bind.conf)

  ```
  $JBOSS/server/bin
  ```

- if you changed the logging levels (log4j.xml)

  ```
  $JBOSS/server/default/conf
  ```

- if you added some .jar files for specific JDBC drivers

  ```
  $JBOSS/server/default/lib
  ```

- if you changed JBoss configuration (like mail-service.xml)

```
$JBOSS/server/default/deploy
```

44.2. Backup before an upgrade

If you plan an upgrade, you may backup separately the configuration files in order to easily apply again your configuration on the default one (take care not to loose any evolution on these files):

- Main configuration is in `$JBOSS/server/default/deploy/nuxeo.ear/config`
- Datasources are defined in `$JBOSS/server/default/deploy/nuxeo.ear/datasources`
- Compass backend is configured in `$JBOSS/server/default/deploy/nuxeo.ear/system/nuxeo-platform-search-compass-plugin*/compass.cfg.xm`

44.3. Restore

All you need is to restore previously saved database(s), files and directories.

44.4. Reset

You can simply reset Nuxeo by removing its data: delete `$JBOSS/server/default/data/` and empty all used databases. On start, Nuxeo EP recreates all its tables, files and directories.

Chapter 45. The Nuxeo Shell

45.1. Overview

The Nuxeo Shell is a command line tool for everyone who needs simple and quick remote access to a Nuxeo EP server. You can see it as the Swiss army knife for the Nuxeo EP world. It can be used by regular users to browse or modify content, by advanced users to administrate Nuxeo servers, or by programmers to test and debug.

The Nuxeo Shell is designed as an intuitive and extensible command line application. It can both serve as a user, administrator or programmer tool, or as a framework to develop new Nuxeo command line clients.

The application is based on the Nuxeo Runtime framework and thus offers the same extension point mechanism you can find on the server application.

The main features of the command line tool are:

- Two operating modes: an interactive and a batch mode.
- Advanced command line editing like:
  - auto-completion
  - command history
  - command line colors
  - command line shortcuts like: CTRL+K, CTR+A, etc.
- JSR223 scripting integration. You can thus connect and execute commands on a Nuxeo server in pure scripting mode.
- Extensibility - using Nuxeo Runtime extension points

The advanced command line handling is done using the JLine library.

The only requirement is Java 5+. The tool works on any Unix-like OS, on Windows and on Mac OS X.

45.2. User Manual

The Nuxeo Shell application is packaged as a zip archive. To install it, you need to unzip and copy the content in a folder.

The application folder structure is as follow:

```
+ nxshell
  + app
    + bundles
    + config
    + data
    + lib
  + lib
  - Launcher.class
  - log4j.properties
  - launcher.properties
  - nxshell.sh
```

- The `lib` folder contains JARs needed by the application launcher. The `Launcher.class` is the application launcher and `launcher.properties` contains configuration for the launcher.
- `log4j.properties` is as you expect the configuration for log4j.

- `nxshell.sh` is a shell script that launches the application.
- The `app` folder contains the application code and data.
  - `app/bundles` contains the application OSGi bundles. These bundles will be deployed using a minimal implementation of OSGi (the same that is used on the server side). We may replace the default implementation by equinox later.
  - `app/lib` contains third party libraries needed by the application bundles.
  - `app/config` contains the application configuration files.
  - `app/data` contains temporary data.

The only file you may need to change is the `nxshell.sh` script. Here is the content of this file:

```
#!/bin/bash

#JAVA_OPTS="$JAVA_OPTS -Xdebug -Xrunjdwp:transport=dt_socket,address=127.0.0.1:8788,server=y,suspend=y"
JAVA_OPTS="$JAVA_OPTS -Dorg.nuxeo.runtime.1.3.3.streaming.port=3233"
java $JAVA_OPTS Launcher launcher.properties $@
```

If you uncomment the first line, you will be able to run Nuxeo Shell in debug mode.

The second line [must] be commented out if on the server you are running on a nuxeo runtime 1.4.SNAPSHOT. When running against Nuxeo EP 5.1 you need to let this uncommented.

You can run the application in two modes:

1. In batch mode - in this mode you specify the command that will be executed. After the command execution the process will exit. The command has to be passed as the first argument.

   Here is an example of a command executed in batch mode:

```
$ ./nxclient.sh export /path/to/remote/doc /path/to/local/folder
```

   Here is an example of a command executed in batch mode while also passing parameters:

```
$ ./nxclient.sh ls --host 192.168.2.54
```

2. In interactive mode - in this mode you can share a single session to run multiple commands against the remote Nuxeo EP. To start the application in that mode you should run the **interactive** command in one of two ways:

```
$ ./nxclient.sh interactive
```

```
$ ./nxshell.sh
```

   After entering in interactive mode a command prompt will be displayed. At this prompt you can enter commands in the same way you specify them on the command line in batch mode.

   When not connected to a server, the prompt will be displayed as:

```
|>
```

   After connecting to a server named, let's say "nuxeo-platform", the prompt will be

displayed as:

```
|nuxeo-platform>
```

So, as we've seen in the above examples, executing a command is done by using the following syntax:

```
command [options] [parameters]
```

where options and parameters are optional and are specific to the command you run.

Example:

```
import -u /path/to/doc /path/to/local_file
```

In this case "import" is the command, "-u" is a flag option and "path/to/doc" and "/path/to/local_file" are both parameters

Parameters are stand alone arguments (they are not bound to a specific option) and can be retrieved programmatically using their index. In the example above, the first parameter will have the index 0 while the second the index 1.

45.2.1. Command Options

Command options are defined by a name and optionally a shortcut (a short name). When referring to an option using its name you should prefix it by two hyphens '--'. When using short names you should only use one hyphen as a prefix. For example if you have an option named "host" and using a short name of "h" the following commands are equivalent:

```
$ ./nxshell.sh interactive -h localhost
$ ./nxshell.sh interactive --host localhost
```

Options may take values or may be used as flags turning on / off a specific option by their simple presence.

When using long names you should specify the option values immediately after the option. However when using short names you can group options together. Let say for example we have a command that support 4 options: a, v, i, o. a and v are flag options and both i and o takes an argument as value. In this case the you can group options like the following:

```
command -avi in_file -o out_file
```

or

```
command -avio in_file out_file
```

or anyhow you want. You should keep in mind that when grouping options that take arguments, these arguments will be assigned to each of this options in the order they were specified on the command line.

45.2.1.1. Global Options

Besides the specific options that commands may define, there are several global options that apply to all commands. These are:

- host (**--host** or **-h**) the Nuxeo EP host where to connect - defaults to localhost

- port (**--port** or **-p**) the Nuxeo EP port where to connect - defaults to 62474
- username (**--username** or **-u**) the username to use - defaults to the "system" user
- password (**--password** or **-P**) the password to use
- debug (**--debug** or **-d**) to start with debug mode (logs at DEBUG level)

45.2.2. Commands

There is the list of all built-in commands of nuxeo shell.

Notes:

1. At the time of this writing some of these built-in commands are not yet implemented. Those are be marked with an asterisk (*).
2. The commands can be grouped in two categories:
   - offline commands - command that doesn't need a connection to work. Example: help, log etc.
   - online commands - commands that requires a connection to work. These commands are automatically connecting if no connection is currently established.
3. Some commands make no sense and are not available in both modes (batch and interactive). This will be specified by each command if it is the case.

45.2.2.1. interactive

Runs in the interactive mode. This command is not available when already in interactive mode.

Has no specific options.

```
$ ./nxshell.sh interactive
```

45.2.2.2. help

Displays the help page.

Takes an optional parameter which is the name of a command.

By default, displays the global help page. If a command is given as an argument, displays the help page for that command.

```
$ ./nxshell.sh help ls
```

45.2.2.3. log

Manage logs from console without having to edit log4j.xml. Allow to switch on/off a debug mode, add/remove loggers, ...

log filename [log level [package or class]] creates a file (created in $PWD/log/ if **filename** do not contain path separator, else given path is used). It will record logs at level greater or equal to **level** (if specified, else default level is INFO; available values are TRACE, DEBUG, INFO, WARN, ERROR, FATAL). If a **package or a class** (canonical name) is specified, logs will be filtered from this package/class.

log off [package or class] stops logging. Applies on all custom appenders (console and filenames) if no **package or class** is specified. Don't worry about the warning **log off** command may cause (the logger is re-created):

```
log4j:WARN No appenders could be found for logger (org.nuxeo.ecm.shell.commands.InteractiveCommand).
log4j:WARN Please initialize the log4j system properly.
```

log debug switches debug mode on/off. Same result as starting the shell with **-d** (**--debug**) option. This decrease the log level to DEBUG on console and filenames.

See log4j documentation for more information.

Available only in *interactive* mode

45.2.2.4. connect *

Connects to a given server. If a connection is already established, close it first.

Available only in *interactive* mode

45.2.2.5. disconnect *

Disconnects from the current connected server. If no connection exists, does nothing.

Available only in *interactive* mode

45.2.2.6. ls

Lists the children of the current folder in the repository.

By default, Folderish types are displayed in blue.

Note that the **ls directory-name** command is accepted but won't list the content of `directory-name`, it will only list the content of the current directory. This might be improved in the future to provide a behavior alike the one of the Unix ls command.

Available only in *interactive* mode

45.2.2.7. tree

Displays the complete tree structure of the documents as it would be returned by the Unix **tree** command.

Available only in *interactive* mode

45.2.2.8. cd

Changes the current directory in the repository (to a Folderish document)

45.2.2.9. pwd

Displays the current path in the repository.

45.2.2.10. view

Views info about document. The information displayed can be controlled using these command options:

--all (-a) - view all data

--system (-s) - view only the system data

--acp - view the ACLs on that document

45.2.2.11. rm

Removes a document or a tree of documents.

45.2.2.12. mkdir

Creates a Folder document.

45.2.2.13. put

Uploads a blob to a file document. If the target document doesn't exists, creates it.

45.2.2.14. putx

Creates a document other than a file. Metadata (and blobs) are specified in the Nuxeo export format.

45.2.2.15. get

Downloads the blob from a File document.

45.2.2.16. getx

Gets a document as an XML file (as export, but only for a document).

45.2.2.17. export

Exports documents from the repository.

45.2.2.18. import

Imports documents into the repository.

45.2.2.19. script

Makes it possible to run external scripts. Here are examples of such scripts.

From interactive mode use:

```
script --file <path_to_your_script> <args>
```

45.2.2.20. chperm *

Changes a privilege on the given document.

45.2.2.21. useradd

Creates new user(s).

To create user joe:

```
useradd joe
```

To create the users define in the users.csv CSV file:

```
useradd --file users.csv
```

45.2.2.22. groupadd *

Creates new group(s).

45.2.2.23. groupmod

Modifies a group

To add user joe to the company-service1 group:

```
groupmod --user joe company-service1
```

To set the users of the company-service1 group to be only joe

```
groupmod --set --user madarche company-service1
```

To add the users in the users_for_group.csv CSV file to the company-service1 group:

```
groupmod --file users_for_group.csv company-service1
```

To set the users of the company-service1 group to be the users define in the users_for_group.csv CSV file:

```
groupmod --file users_for_group.csv company-service1
```

45.2.2.24. select

Search the repository using the NXQL query language. For example:

```
select * from Document where ...
```

This command can be used as a starting point to write another script to perform search service queries if needed.

This command is only available in the Nuxeo 5.2 branch.

45.2.2.25. lstypes *

45.2.2.26. viewtype *

45.2.2.27. lsusers *

45.2.2.28. lsgroups *

45.2.2.29. viewuser *

45.2.2.30. viewgroup *

45.3. Troubleshooting

45.3.1. Check listened IP

Be sure on what precise IP address and what precise port the server is listening to. To find out, use the **netstat** command on the server host. The example below shows that, on the same host, there are 2 servers listening on the 62474 port but on different IP addresses. You should always use as arguments to the nuxeo shell the exact IP address and port you find out.

```
$ sudo netstat -ntlap | grep 62474
tcp        0      0 ::ffff:192.0.0.11:62474     :::*                        LISTEN      3623/java
tcp        0      0 ::ffff:192.0.0.10:62474     :::*                        LISTEN      1346/java
```

Alternatively you can use the **lsof** command.

```
$ sudo lsof -i :62474
COMMAND   PID       USER   FD   TYPE    DEVICE SIZE/OFF NODE NAME
java     3623      jboss  266u  IPv6 0x5d69a70      0t0  TCP [::192.0.0.11]:62474 (LISTEN)
java     1346      jboss  266u  IPv6 0x5d7af30      0t0  TCP [::192.0.0.10]:62474 (LISTEN)
```

45.3.2. Check connected server

Be sure to be connected on the right server. To do so, issue a **view** command. This will display information on the current context:

```
192.0.0.10> view
------------------------------------------------------------------------
UID: 2a13db70-f133-473c-9a90-6838d01610aa
Path: /
Type: Root
------------------------------------------------------------------------
Title: 2a13db70-f133-473c-9a90-6838d01610aa
Author: null
Created: null
Last Modified: null
------------------------------------------------------------------------
Description: 2a13db70-f133-473c-9a90-6838d01610aa
------------------------------------------------------------------------
```

45.3.3. Multi-machine case

When working with multi-machine installation, connect the Nuxeo Shell on the server running Nuxeo Core (i.e. the stateful server on a bi-machine configuration).

45.4. Extending the shell

If you a need a new command (a new functionality) not provided by the Nuxeo Shell, you can add it very simply by writing a Nuxeo plugin in exactly the same manner you would for a Nuxeo server instance. This is done by writing a Java class for each new shell command and declaring (that is registering) each command in a single XML file as a contribution to an extension point.

Here is an example of how to register two imaginary new custom commands **addapplicants** and **purgeobsoletedocs**.

45.4.1. Registering New Custom Commands

This is how to register the new commands:

```
<?xml version="1.0"?>
<component name="com.company.nuxeo.shellcommands" version="1.0">

  <documentation>
    Extra Nuxeo Shell commands shown as examples for the Nuxeo Book.
  </documentation>

  <extension target="org.nuxeo.ecm.shell.CommandLineService" point="commands">

    <command name="purgeobsoletedocs" class="com.company.nuxeo.shellcommands.ObsoleteDocsPurgeCommand">
      <description>
        Purge obsolete documents considering based on predefined hard-coded logic.
      </description>
      <help>Purge obsolete documents</help>
    </command>

    <command name="addapplicants" class="com.company.nuxeo.shellcommands.ApplicantsAdderCommand">
      <description>
        Adds applicants by creating the right folders with the right permissions.
      </description>
      <params>
        <param index="0" type="path" />
      </params>
      <help>Adds applicants</help>
    </command>

  </extension>

</component>
```

45.4.2. Java Code for the new commands

This is how to write a Java class for a new command:

```


public class ApplicantsAdderCommand extends AbstractCommand {

    public static final String COMMAND_NAME = "addApplicants";

    public void run(CommandLine cmdLine) throws Exception {
        String[] elements = cmdLine.getParameters();
        // parse command line
        if (elements.length != 1) {
            log.error("Usage: " + COMMAND_NAME + " file");
            return;
        }
        File file = new File(elements[0]);

        addApplicants([...]);
    }

    void addApplicants([...]) {
    }

}
```

45.4.3. Building the shell plugin

Finally the `pom.xml` file for the plugin to the Nuxeo shell needs to at least contain the following dependency to be able to build extend the `AbstractCommand` class. Building the plugin will generate a `XXX.jar` file. Note that you should of course replace the *version* given as an example with the version suited for your need.

```
<dependencies>

  [...]
```

```
  <dependency>
    <groupId>org.nuxeo.ecm.platform</groupId>
    <artifactId>nuxeo-shell-commands-base</artifactId>
    <version>5.1.7-SNAPSHOT</version>
  </dependency>

  [...]

</dependencies>
```

45.4.4. Deploying the shell plugin

1. Install or decide on using a nuxeo shell already installed program.
2. Copy the generated `XXX.jar` file into the `bundles` directory of the nuxeo shell installed program.
3. Delete the `data` directory, if there is any, of the nuxeo installed program. This is to purge any caching of registered JARs.

The next time the nuxeo shell is restarted your new commands will be available.

Part VI. Core developer guide

Chapter 46. Coding and Design Guidelines

46.1. Introduction

> "Follow a common coding standard, so that all the code in the system looks as if it was written by a single — very competent — individual. The specifics of the standard are not important: what is important is that all the code looks familiar, in support of collective ownership."
>
> —Ron Jeffries *XProgramming.com*

> Coding Standards are a good idea. Every team should adopt a coding style and standard, and stick to it. The code produced by that team should have a consistent look and feel that is devoid of individual preferences and fetishes.
>
> —"Uncle" Bob Martin *ObjectMentor website*

The primary goal of this chapter is to provide common conventions, as well as guidelines that aim at ensuring a high level of quality (with respect to maintainability, extensibility, modularity, security, testability, documentation, etc.) throughout the Nuxeo project.

As such, it is primarily written for the platform's developers, both "core" developers and contributors, and developers that write extension modules that will someday find a place in the Nuxeo codebase.

If you're working on your own project with no plan to contribute code to the Nuxeo platform, you will probably be still interested in reading this chapter, and adapt the conventions stated here to your company's or organization's own coding standards.

46.2. External Coding Standards

Rewriting a whole coding standard for the Nuxeo project would be a poor use of our time, since there are already several documents and books that do a fine job in this respect.

We've opted for a more pragmatic, two-pronged approach:

1. reference reputable external guides or books, and document any difference
2. propose default settings for common Eclipse tools (including the built-in code formatter and the CheckStyle plugin), tuned to comply to these conventions.

The official coding standard for the project is: "The Elements of Java Style" [Vermeulen2000] which is a little book that can be bought from Amazon or other book dealers.

If a PDF suits you better, you can download an earlier version of the book from here.

Note however that these guidelines have been written in 2000 (last century!) and some of the recommendations need the be updated in light of this millenium's experience.

There are also architectural guidelines that need to be followed. At this moment, they are written here (section 4 - "Guidelines")

46.3. Some points that need attention

46.3.1. Java code formating

> Readability counts.
>
> —Tim Peters *The Zen of Python*

Java code should be formatted using Sun's Code Conventions for the Java Programming Language.

Regarding code formatting (which is just one small part of the book and PDF mentioned above), we'll be using the standard Eclipse default style, with a couple of minor tweaks (see this README.txt for how to configure Eclipse to support this style).

The major points are:

- 4 spaces (no tabs!!!) indentation
- spaces before and after = signs and most binary operators
- spaces after ,
- no space after *(* or before *)*
- regarding code block, we are using 1TBS ("One True Brace") style:

  Bad:

```
if (xxx)
{
    yyy
}
```

  Good:

```
if (xxx) {
    yyy
}
```

- make a block even for only one statement:

  Bad:

```
if someThing() doSomething;

if someThing()
    doSomething;
```

  Good:

```
if someTest() {
    doSomething;
}
```

- Don't prefix your instance variables by "this." when it's not mandatory.

  Why? Because with modern IDEs, instance variables are typeset in a different color than local variables, hence you don't need the extra information provided by the "this." prefix to recognize which is which.)

- etc.

46.3.2. XML code formatting

1. XML code should also be formatted, using 2 spaces for each indent level (not 4 spaces).

   Badly formatted XML code can be reformatted using the Eclipse formatter.

   Always check that the result is better than the original, and tweak afterwards if needed.

46.3.3. Design

1. Use interfaces every time it is possible instead of concrete classes.

   Ex: declare a variable or argument as a "Map", not as a "HashMap".

   Why? Using interfaces makes it possible to use several different implementations. Actually, interfaces are the basis of component-oriented programming.

2. But don't overdo it if it's not necessary.

   Why? The more code we have, the more expensive maintainance is.

3. Avoid cyclic dependencies.

   Why? This is a common design principle, called the "Acyclic Dependency Principle" or "ADP".

   How? use JDepend or a JDepend plugin for you IDE (ex: http://andrei.gmxhome.de/jdepend4eclipse/) or look at the JDepend report by Maven.

   More info:

   - How JDepend Changed My Java Packaging
   - Managing Your Dependencies with JDepend
   - Uncle Bob Martin's "Design Principles and Design Patterns" (pages 18-22, "The Acyclic Dependencies Principle (ADP)")

46.3.4. Unit tests

> If it's not tested, it's broken.
>
> —Bruce Eckel

1. Write unit tests. A lot.

   How? Learn how to use JUnit. Use tests in existing modules as examples. Remember that well-tested modules have at least as many lines of test code than not test code.

2. Check that your unit tests have a coverage of most of the critical parts of your code.

   How? Look at the Cobertura report by Maven, or use the TPTP or EMMA (http://www.eclemma.org/) plugins for Eclipse or use the native code-coverage function of IDEA.

3. Design your API so as to make unit testing possible and easy.

   See API Design As If Unit Testing Mattered

Some cosmetic remarks related to writing unit tests:

1. Use "assertEquals(constant, value)", not "assertEquals(value, constant)"

   Why? Because one has to choose a convention and this one

2. Use "assertNull(value)", not "assertEquals(null, value)"

   Why? Because code is shorter and looks cleaner this way.

46.3.5. Security

1. Read Graff and van Vyk's "Secure Coding" [Graff2003] book. More info on the securecoding.org website.

2. If you don't have access to the book, read at least Secure Coding: The State of the Practice (including the Architectural Principles, the Design Ideas and the Java Tips sections) from one of the two authors.

3. Read these slides from JavaOne 2006.

46.3.6. Naming convention

1. Don't use underscores ("_") in variables, methods or class names, use camelCase notation instead.

   Why? Because it is the Java convention.

2. Check some good articles on naming things, like Ottinger's Rules for Variable and Class Naming.

   It is especially important that we all use the same words (in case there is an ambiguity) to refer to the same concepts (like "fetch" vs. "retrieve", for instance).

46.3.7. Information hiding

1. Minimize the accessibility of classes and members (for instance by making them "private" unless you have a reason to believe that they will be used from other classes, and more generally by using the weakest access modifier), so as to keep you API clean and easy to understand for other developers.

   Reference: "Effective Java" [Bloch2001], item 12.

46.3.8. Use modern Java features

1. Use Java 5's "foreach" construct instead of explicit iterators whenever possible (which is not always the case).

2. Use java 5's generics.

   Why? This will improve the amount of static checks done by the compiler.

46.3.9. Logging

Follow this piece of advice from Alex Miller:

My opinion is that there are 4 log levels that matter for general program logging (specialized logs like audit and performance have their own rules). I dont care what you call them so feel free to replace these with the names of your choice.

- ERROR - an error has occurred in the program (usually an exception). All internal or unexpected failures should be logged at ERROR level. Programmers should care about anything logged as an error. User errors (like input validation) should generally NOT be logged at ERROR level. These are expected (if invalid) inputs and you should respond to the user appropriately (another post altogether), but this is not a program error. Exceptions may be made for things like repeated invalid password attempts that are likely security problems.

- WARNING - an anomalous condition has been detected and the program will attempt to deal with it. No action needs to be taken externally but you might be interested if things go down in flames 5 minutes later. The conditions for warnings are relatively rare but an example would be reaching a warning threshold on a connection pool or a loss of connectivity that can be repaired by reconnecting to the same or different source.

- INFO - an interesting piece of information that helps to give context to a log, often when things are starting or stopping. INFO messages should never be printed on a per transaction (whatever that means to you) path through the code. That is, INFO messages should be 1-time or very occasional messages.
- DEBUG - whatever your OCD programmer mind desires.

The first three levels (ERROR, WARNING, and INFO) should always be on. If your logs grow rapidly with these levels on, something is wrong. You should be able to run in production with these levels for weeks without a problem (but you should really use a rolling log writer just in case). DEBUG should be used only during development for debugging.

46.3.10. Documentation: Comments and Javadoc

> Reuse is something that is far easier to say than to do. Doing it requires both good design and very good documentation. Even when we see good design, which is still infrequently, we won't see the components reused without good documentation.
>
> —D. L. Parnas *The Mythical Man-Month: Essays on Software Engineering*

1. Always use comments to document what is not obvious.

   Why? Otherwise, it will be harder for others (even you!) to maintain or extend your code. Some people may even end up refactoring your code under wrong assumptions.

   See for instance item #28 of Bloch's "Effective Java" [Bloch2001] for more insight, including this most important principle:

   > The doc comment for a method should describe succinctly the contract between the method and its client. With the exception of methods in classes designed for inheritance, the contract should say what the method does rather than how it does its job.

2. Write javadocs according to the javadoc manual and the official Sun javadocs guidelines.

   Why? Because javadocs are compiled by a standard tool, you must conform to the syntax mandated by this tool.

3. Common javadoc guidelines that are often forgotten include:
   - Don't include XML tags without escaping the < and > characters.
   - Remember that the first sentence (i.e. everything until the first dot) is used as a short description of the package / class / method / etc. it documents. Hence end all your sentences with a dot.
   - Use for separating paragraphs, not <br> or
   - Don't use in the source though as it is useless and rather bad looking.
   - If your javadoc starts with "This class" or "This methods", this is most probably wrong.
4. Proofread carefully (including checks for spelling and grammar errors) and check that your javadocs provide useful information for third-party developers (or your teammates) by actually reading the generated doc.

   How? See either "Project -> Generate Javadocs" from Eclipse, or "mvn javadoc:javadoc" from the project root and browse the docs.

   Remember that you can see a preview of how your javadoc will look like from Eclipse by using the "Javadoc" tab in the Java perspective.

5. Don't mix Javadoc comments with implementation comments.

Javadocs comments are intended to be read by clients of the API, implementation comments by people who need to review you code or who will end up maintaining it (including yourself).

For instance, don't put "TODO" markers in your javadoc because they are intended for yourself or other people co-developing or maintaining your code, not clients of the API.

6. Start by documenting the most important and less obvious things.

   If you don't have enough time to document everything, document first the API part of your code (usually, the interfaces), because that's what third-party developers are going to use.

   Also documenting the non-obvious (like giving an overview of a package or an important class) is more important than writing for instance, that a "getTitle()" method "@returns the title".

7. Write package-level javadocs.

   Package-level Javadoc (just `package.html` files in your source code) are the most important javadocs you have to write, because that's what people reading the javadoc-generated site will read first, and because that's where the information is usually the less obvious.

8. Sign your code (in the modules headers).

   It is very important that people who will read / maintain your code know that you are the author, so that they can ask you questions, give you feedback, etc.

9. When you've borrowed code from another open source project, always document it so that:

   - we are sure that there is no license conflict with the original code
   - we understand why the code in question doesn't follow our own coding conventions or isn't unit-tested as it should (it should anyway, but this is another story)

10. Put markers as (inline) comments to document parts of your code that will need further work.

    - use FIXME for really serious (like, release-blocking) issues.
    - use the TODO for remaining tasks.
    - don't use XXX which is usually associated with inappropriate content.
    - use BBB to mark code used to ensure compatibility with a deprecated feature, that will be removed after a certain release.
    - do not leave TODO markers behind when they are not relevant (for instance for dummy classes used in tests and auto-generated by the IDE from an interface or an abstract class).

11. Look critically at the javadoc-generated site and try to improve it.

    Either go to http://maven.nuxeo.com/apidocs/ and check the apidoc for your project, or run *mvn javadoc:javadoc* locally and browse the generated apidoc, and ask yourself the simple question: "if I was a third-party developer, would I understand how to use the API by reading this".

46.3.11. Deprecation

1. Don't use deprecated features (= API), either from Java or third-party libraries, or from the Nuxeo framework.

   Hint: they should appear as stroked-out in your IDE.

2. Deprecate your own API when you have to, but make sure you write comments that explain to your API's clients how to migrate from the old API to the new one. Use BBB markers (see above) if needed.

3. Deprecated APIs should be maintained at least for 1 release cycle, at least for external clients of the APIs, but we should strive to switch to the new APIs internally in 1 release cycle.

46.4. Methodology tips

Here are a few points and tips to keep in mind.

46.4.1. Use the power of your IDE (and its plugins)

Modern IDEs (+ adequate plugins, if necessary) include many code-checking functions that help find out issues:

1. For Eclipse, the simplest plugin to use is TPTP.

   How?

   - Use the update manager to get the TPTP plugins.
   - Right-click on your project, and select "Analysis".
   - Enable all the rules, and run the analysis on your module.
   - Fix the issues that appear serious (you still have to think).

2. You should also use the Checkstyle and FindBugs Eclipse plugins to ensure minimal bug count and maximal coding style coherence.

   See these great slides from JavaOne 2007 for instance.

3. Read Improving code with Eclipse plugins on developerWorks for more background information on the subject.

   From the article:

   > This article covers what I consider to be the "big five" code analysis areas:
   >
   > - Coding standards
   > - Code duplication
   > - Code coverage
   > - Dependency analysis
   > - Complexity monitoring
   >
   > These analysis areas can be uncovered using a number of the following slick Eclipse plugins:
   >
   > - CheckStyle: For coding standards
   > - PMD's CPD: Enables discovering code duplication
   > - Coverlipse: Measures code coverage
   > - JDepend: Provides dependency analysis
   > - Eclipse Metrics plugin: Effectively spots complexity

   NB: Coverlipse may or may not work correctly, an alternative is EclEmma which is very similar.

46.4.2. Refactor

1. When something looks wrong in the code ("smell"), refactor it, but make sure you don't break anything.

   How?

   - Check that the code you are refactoring is covered by some unit tests.
   - Refactor. (Tip: a modern IDE (Eclipse and IDEA for instance) has some functions that may help.)
   - Check that the code still compiles (including dependent modules) and all the unit tests are still passing, and commit.

46.5. Important references

Here is a list of useful stuff to read:

1. Joshua Bloch's "Effective Java" [Bloch2001] book. Highly recommended. Probably the best "advanced" Java book.
2. Joshua Bloch's "Designing Effective API" slides and video.
3. ObjectMentor's design articles (click on "Design Principles"), the most important one being Principles and Patterns.

Chapter 47. Development Tools and Process

This chapter describes useful development tools.

47.1. Code Quality with Eclipse Plugins

Eclipse can benefit for several plugins for improving code quality, in both adherence to the project's coding standard and in removing bugs.

47.1.1. Using Checkstyle

TODO

47.1.1.1. The Short Story

If you are already familiar with the Checkstyle Eclipse plugin, just configure it to use the `checkstyle.xml` at the root of the SVN directory for nuxeo-ep:
`http://svn.nuxeo.org/trac/nuxeo/browser/nuxeo-ep/trunk/checkstyle.xml`

47.1.1.2. Longer Story: installation and Configuration

TODO

47.1.1.3. Notes on Reported Issues

TODO

47.1.2. Using TPTP

TODO

47.1.3. Using FindBugs

TODO

47.2. Profiling with NetBeans Profiler

IDEA doesn't provide an integrated profiler and Eclipse's profiler, provided by the TPTP project, doesn't currently work on Mac OS, one can use NetBeans to profile the Nuxeo platform.

Here is how to do it:

1.

47.3. NXPointDoc Documentation tool

Nuxeo heavily uses extension points. In order to manage them nxPointDoc tool has been created. Its purpose is to explore all XML files and build the documentation of each explored components. Cross links and indexes are also built to ease the navigation.

The NxPointDoc pages are available at `http://svn.nuxeo.org/nxpointdoc/`

The tool is written in Python and the following libraries are needed:

- Genshi for templating
- ElementTree for XML processing
- Pygments for code highlighting

47.3.1. Documenting a component

A component XML file is structured as follow:

```
<component ...>

  <!--############### Component configuration ############-->
  <!-- implementation class (optional) -->
  <implementation>...</implementation>

  <!-- component properties (optional) -->
  <property name="..." value="..."/>
  ...
  <property name="..." value="..."/>

  <!--############### Extension points ################-->
  <!-- extension points are optional -->
  <extension-point ...>
      ...
  </extension-point>
  ...

  <!--############### Contributions ################-->
  <!-- contributions are optional -->
  <extension ...>
      ...
  </extension>
  ...

</component>
```

We can see that the only required element is the component element (although it is useless to have an empty component). So there are 3 main sections (any of these sections are optional).

- Component configuration: This section defines the component implementation class and some properties to initialize the component (This section content may be modified in future especially when aligning nuxeo components with OSGi services).
- Extension points: This sections contains all the extension point declared by the component.
- Contributions (`extension` tag): This section contains all the contributions made by this component to other components.

To add documentation to these elements a `<documentation>` tag will be used. An element may have different content depending on what it is documenting. While some information is already available in other XML elements in the file, there is no need to duplicate these information inside the documentation provided though the element. For example the name of the component can be retrieved from the name attribute of the component element, the implementation class name from the implementation element etc.

To format the description text, we can use XHTML tags and javadoc-like markers such as `@property`, `@schema` etc. Javadoc-like links are also supported: `@see` points on the javadoc, `@component` points on another component documentation. For example, `{@see org.nuxeo.ecm.core.schema.types.Type}` will point on the `http://maven.nuxeo.org/apidocs` corresponding page while `{@component org.nuxeo.ecm.webapp.directory.DirectoryTreeService}` will explicitly insert a link to the related NxPointDoc page. Code colorization is also supported through the `<code language='xml'> ... </code>` tags. If no language is given, `xml` is taken by default. Java (`language='java'`) and many other languages are supported (see Pygments pages).

Regarding `<component>` documentation, the following elements are available:

- `@author`: may be duplicated for multiple authoring

- `@version`
- `@property`
- `@see`: points to Javadoc
- `@component`: points to nxpointdoc
- `@deprecated`
- `component@name` attribute: the component name
- `component/implementation` tag: the implementation class
- `component/require` tag: required elements
- `component/documentation` tag: the description

For `<extension-point>` the following elements have to be used:

- `@author`: may be duplicated for multiple authoring
- @schema
- `@deprecated`
- `@see`
- `@component`
- `component/extension-point@name` attribute: the name
- `component/extension-point/documentation` tag - the description
- If the extension point is using `object` sub-elements, the `DTD` should be extracted from the XMap annotated class, otherwise the user may specify the `DTD` using the `@schema` marker inside the `documentation` element

For <extension>, describing contributions to an extension-point, we have

- `@author`: may be duplicated for multiple authoring
- `@see`
- `@component`
- `@deprecated`
- `component/extension@target` attribute - rendered as a link to the component documentation
- `component/extension@point` - rendered as a link to the extension point documentation
- `component/extension-point/documentation` description

Here is a short example of what a component xml file may look like.

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.MyService">
  <documentation>
  My demo service
  <p/>
  This service does nothing
  @property home home directory to be used to create temp files
```

```
  @property timeout the time interval in seconds
  @version 1.0
  @author Bogdan
  </documentation>
  <require>org.nuxeo.ecm.Service1</require>
  <require>org.nuxeo.ecm.Service2</require>
  <implementation class=”org.nuxeo.ecm.core.demo.Service2”/>
  <property name=”home” value=”/home/bstefanescu”/>
  <property name=”interval” value=”20” type=”Integer”/>

  <extension target="org.nuxeo.ecm.SchemaService" point="schema">
      <documentation>Common and Dublin Core schemas</documentation>
      <schema name="dublincore" src="schema/dublincore.xsd" />
      <schema name="common" src="schema/common.xsd" />
  </extension>

  <extension-point name="repository">
      <documentation>Register new repositories</documentation>
      <object class="org.nuxeo.ecm.RepositoryDescriptor"/>
  </extension-point>

</component>
```

Feel free to browse the NxPointDoc site and teh corresponding xml file to go deeper. The systematic link to the source code svn repository may help you.

47.3.2. Creating the NxPointDoc site

nxpointdoc is a command line program that creates the whole site from a source repository `SOURCE_DIR` to a target publication directory `TARGET_DIR`. Each component is analyzed and all related pages created. Index pages are then created.

```
$ ./nxpointdoc.py -h
usage: nxpointdoc.py [options]

options:
--version                         show program's version number and exit
-h, --help                        show this help message and exit
--source=SOURCE_DIR               Source root directory containing xml component files
--target=TARGET_DIR               Target directory for the generated documentation
--template=TEMPLATE               Genshi template for component html file
--template-index=TEMPLATE_INDEX   Genshi template for index html file
--allow-xhtml-comment=ALLOW_XHTML_COMMENT 'no' to not interpret xhtml tags in comment
--color=COLOR_CODE                'no' to not color <code> contents
```

Valid `SOURCE_DIR` and `TARGET_DIR` are mandatory. Template files have to exist. The one delivered have the `.template` extension and can be used as is.

47.3.3. Browsing NxPointDoc

NxPointDoc generates 3 indexes that are the entry points; The documentation is accessible at `http://svn.nuxeo.org/nxpointdoc/` with 3 indexes related to components, extension points and contributions. Each one is an entry point for the documentation. The Indexes give the name and the first line of the documentation. An hyperlink allows to see the detail of the examinated item.

The statistic gives some rough indicators on the documentation coverage, globally or for each component file. The `G.D.C` stands for Global Documentation Coverage while the `I.D.C` stands for Individual Documentation Coverage. They show the ratio between all the information/documentation that is written over all the entries that are considered as mandatory (like author, documentation, etc.). The higher these indicators are, the better it is.

47.4. Quality Assurance with continuous integration

"Continuous Integration is a software development practice where members of a team integrate

> their work frequently, usually each person integrates at least daily - leading to multiple integrations per day. Each integration is verified by an automated build (including test) to detect integration errors as quickly as possible. Many teams find that this approach leads to significantly reduced integration problems and allows a team to develop cohesive software more rapidly."
>
> —Martin Fowler - *Continuous Integration (an introduction)*

See http://en.wikipedia.org/wiki/Continuous_Integration

It's as important to follow quality processes on development as to maintain this quality among time. Nuxeo is involved in such practices that will guarantee or reinforce its products quality.

47.4.1. Rules and means

Nuxeo products and tools are continuously built over time, at each change and against multiple environments. Nuxeo QA team sets and maintain a QA environment applying CI rules and so providing to developers means to check their code quality and being warned in case of any problem.

- Maintain a code repository

  Nuxeo sources repositories hg.nuxeo.org and svn.nuxeo.org are under continuous integration.

  This includes Nuxeo EP, Nuxeo addons, Nuxeo RCP, Nuxeo WebEngine, Nuxeo Books, tools and plugins and of course all our customers' projects.

- Builds are automated

  This is done by Hudson on Nuxeo QA Unit, Functional and Integration tests.

- Every commit on mainline is integrated

  When code is committed, target project is built, as all projects depending on it. The full chain is verified, from build to deployment.

  Mainlines on Nuxeo EP and addons are the main branches in development: 5.1 and 5.2 (resp. 1.4 and 1.5 for associated subtrees). For projects under SubVersion, that means the trunk and, if exists, 5.1 branch.

- Everyone can see the results of the latest build

  Hudson plugins ensure to warn potential responsible(s) of build fail by mail and jabber, so they can react quickly.

  Moreover, every build fail is sent on ECM QA mailing list.

- Make it easy to get the latest deliverables

  Nightly builds are done. Produced artifacts are published on our Maven repositories maven.nuxeo.org. Currently managed with Nexus, our repositories store all released artifacts and recent snapshots.

- Keep the build fast

  Continuous Integration is done on multiple servers, more or less powerful, using slaves in order to distribute the load.

  Thanks to Maven and to Nuxeo modularity, each module is built separately and as a consequence, quickly.

- Test in a clone of the production environment

  We have two integration levels: unit and functional.

  First level checks code compilation and runs Unit tests. A lot of Unit tests simulate target environments

(with mock objects). Dependent projects/modules are then added to the CI chain.

Second level runs packaging tools and automated deployment against multiple environments (we aim at covering JVM versions, SQL backends, OS, browsers, performance, ...). Finally we use Selenium tests to check functional integrity. This also indirectly provides a continuous integration on our tools (packaging, convenient scripts, ...).

47.4.2. Quality directives for Nuxeo developers

These practices apply on every script, project or module. They should be strictly followed.

- Code must be under continuous integration.

  Except for prototype and spike solutions (sandbox projects or temporary branches), all projects must be under CI. If not, ask for it to the QA team, providing the informations mentioned in the following Hudson part.

- Automate the build

  Think about QA tools that will have to test the project without any human intervention. Provide Maven, Ant or, in the worse case, Shell autonomous configuration.

- Make your build self-testing

  Think "test-driven development". Simply building a project/module and running its Unit tests should be a valuable measurement of the code stability. Unit tests code coverage often needs to be increased.

- Commit every day

  Smaller are the commits, lower is the risk of conflicting changes and easier is the bug analysis.

- Stay tuned

  Be aware of CI builds, particularly failed builds.

  Log on http://qa.nuxeo.org/hudson/ and check your profile's informations, especially your jabber address. Hudson will then be able to contact you via Jabber when you are suspected of having broken something.

  Subscribe to ECM QA mailing list. Use mail filters to quickly catch and fix problems. Hudson will send you a mail if it detects one of your commits between succeed and failed tests.

  If you're used to, RSS feeds are also available.

  Check regularly your projects health on our QA sites. Inform QA team if you notice any issue.

  Always consider a build failed as an emergency.

47.4.2.1. Maven parent POM

Maven Parent POM file gives a lot of useful information. Take care to fill in you project's `pom.xml` file:

- main tags

```
<name>Nuxeo ECM Projects</name>
<description>Nuxeo ECM Platform and related components</description>
<organization>
  <name>Nuxeo SAS</name>
  <url>http://www.nuxeo.com/</url>
</organization>

<licenses>
  <license>
    <name>GNU LESSER GENERAL PUBLIC LICENSE, Version 2.1</name>
    <url>http://www.gnu.org/copyleft/lesser.txt</url>
```

```
  </license>
</licenses>

<mailingLists>
  <mailingList>
    <name>Nuxeo ECM list</name>
    <subscribe>http://lists.nuxeo.com/mailman/listinfo/ECM</subscribe>
    <unsubscribe>http://lists.nuxeo.com/mailman/listinfo/ECM</unsubscribe>
    <archive>http://lists.nuxeo.com/pipermail/ecm/</archive>
  </mailingList>
</mailingLists>

<issueManagement>
  <system>jira</system>
  <url>http://jira.nuxeo.org/browse/NXP</url>
</issueManagement>

<ciManagement>
  <system>Hudson</system>
  <url>http://qa.nuxeo.org/hudson/</url>
</ciManagement>
```

- "scm" tag

```
<scm>
  <connection>scm:hg:http://hg.nuxeo.org/addons/nuxeo-samples</connection>
  <developerConnection>scm:hg:https://hg.nuxeo.org/addons/nuxeo-samples</developerConnection>
  <url>http://trac.nuxeo.org/nuxeo/browser/nuxeo_samples</url>
</scm>
```

- "developers" tag (there's no rule for tags within the "developer" tag, feel free to add useful information such as "role", "url", "organization" or "module")

```
<developers>
  <developer>
    <name>John Doe</name>
    <email>jdoe@nuxeo.com</email>
  </developer>
</developers>
```

- You also have to add <repositories> section in the project's parent POM in order to make your project fully autonomous.

```
<repositories>
  <repository>
    <id>public</id>
    <url>http://maven.nuxeo.org/public</url>
    <snapshots>
      <enabled>false</enabled>
    </snapshots>
  </repository>
  <repository>
    <id>public-snapshot</id>
    <url>http://maven.nuxeo.org/public-snapshot</url>
    <snapshots>
      <enabled>true</enabled>
    </snapshots>
  </repository>
</repositories>
```

All these tags are intelligently inherited so that, if you're project's Maven parent is nuxeo-ecm or one of its children, you don't have to repeat informations such as "organization", "licenses", "mailingLists", "issueManagement". Also, when working on a project with sub-modules, it's only necessary to set "scm" on the parent POM.

47.4.2.2. Hudson configuration

Adding a project under continuous integration in Hudson requires:

- SCM URL
- Build command

In case of Maven, it's the goals to run (usually, it will be "clean install"). Consider using Maven "Profiles" to manage different behaviors like development versus production environment.

In case of Ant, you may need to provide some parameters on the command line (equivalent to what can be set by a human user in a `build.properties`).

In case of Shell (avoid it as much as possible), it's a simple command with working default values in case of required parameters. If needed, some environment constants may be set.

- Notification target(s): eg, the team mailing list.

47.5. Release process

47.5.1. Overview

Release process is managed and tested by multiple tools:

- Hudson continuous integration which generates a candidate release every night based on snapshots.
- nx-builder
- Selenium tests
- Funkload tests
- IZPack for jar packaging
- Wine and Inno Setup Script for Windows packaging

When a release is wanted, all continuous integration chain must be satisfied, nightly builds are manually tested to complete automated tests and candidate release is "promoted" to public release (code is tagged, artifacts are uploaded to the maven repository and packages are published on Nuxeo web site.

47.5.2. Continuous integration coverage

For now, continuous integration covers those configurations:

- Nuxeo DM.
- Linux Ubuntu (Debian).
- Sun Java 5, Sun Java 6.
- JBoss application server.
- VCS backend on H2.
- VCS backend on PostgreSQL.

Not automatically tested at integration level (full deployement and tests):

- Nuxeo EP.
- Nuxeo-shell and scripts.
- JCR backend with various databases: H2, Derby, PostgreSQL, MySQL, Oracle, ...
- VCS backend with various databases: Derby, MySQL, Oracle, ...

- Various application servers: Jetty, GlassFish (GF3) and Tomcat.
- Various Windows OS.
- Various Java providers.

47.5.3. Help testing release candidates

You can download nighlty candidate releases from IT-nuxeo-5.2-build, test and send feedback on our ECM mailing list or, in case of bugs confirmed, in our Issue Tracker (Jira).

Chapter 48. Packaging Nuxeo EAR

48.1. Introduction

Goal of this chapter is to help creating projects over Nuxeo. How to create a project structure including a packaging module that builds one or more Enterprise ARchive (EAR), ready for deployment. Ant scripts, Maven profiles and Nuxeo assemblies allow defining simple commands to build packages that will fit each target environments.

48.2. Basic project structure

Good practice is to define a parent pom at project root and set a module-ear. Beside that, you may have as many modules and sub-modules as you want. We recommend to split your code according to some development and deployment criteria: API, web or core contributions, stateful or stateless services, ... Following this main rule, you have to create multiple modules where to dispatch your code and resources. Maven will manage the dependencies between modules and our assembly files will package EAR with all required libraries.

Following multi-module project template is available here

```
|-- README.txt
|-- build.properties.sample                                                    ❶
|-- build.xml                                                                  ❷
|-- fixeclipse                                                                 ❸
|-- pom.xml                                                                    ❹
|-- .project.ok                                                                ❺
|-- template-api                                                               ❻
|-- template-core                                                              ❼
|-- template-web                                                               ❽
|-- template-ear                                                               ❾
|   |-- pom.xml
|   `-- src
|       `-- main
|           |-- assemble                                                       ❿
|           |   |-- template-dev.xml
|           |   |-- template-ldap-pg.xml
|           |   |-- template-stateful.xml
|           |   `-- template-stateless.xml
|           |-- resources_template_common                                      (11)
|           |   `-- README.txt
|           |-- resources_template_dev
|           |   `-- README.txt
|           |-- resources_template_ldap_pg
|           |   |-- config
|           |   |   |-- default-repository-config.xml
|           |   |   |-- jdbc-compass-connection-config.xml
|           |   |   `-- sql.properties
|           |   |-- datasources
|           |   |   `-- unified-nuxeo-ds.xml
|           |   `-- system
|           |       `-- nuxeo-platform-search-compass-plugin-5.1.7-SNAPSHOT.jar
|           |           `-- compass.cfg.xml
|           |-- resources_template_stateful
|           `-- resources_template_stateless
|-- template-stateful-services                                                 (12)
|-- template-stateless-services                                                (13)
`-- template-test                                                              (14)
```

❶ Ant sample properties. Copy this file to `build.properties` and personalize it to fit your installation.
❷ Ant script. It's aimed to ease the use of maven, giving targets to automatize common tasks on project (test, compile, deploy, package, ...).
❸ OS dependent, will only work on Linux/Unix and Mac OS X. Copy it from Nuxeo EP root. It's a shell utility script that calls maven-eclipse-plugin to create or update Eclipse's `.project` and `.classpath` files, taking care of setting different build directories for Maven and Eclipse. The script gives precedence to any present `.project.ok` or `.classpath.ok`.
❹ Maven project parent file.
❺ Required to import this project into eclipse as a whole resources project (each module will also be individually imported as a java project).
❻ This module contains common code usable in all other modules, such as interfaces, adaptors, constants,

data transfer objects, ...
❼ This module contains all contributions to Nuxeo Core.
❽ This module contains all JSF, Seam, WebEngine contributions, components and templates.
❾ This module provides packaging of Nuxeo EP following assembly descriptors.
❿ Assembly descriptors. Here are some samples.
Resources directory sample.
This module gathers all Message Resolved Beans that modify some data and all services that define their own persistence source.
This module contains all services that are not going into template-stateful-services.
This module is for use in unit tests from other modules.

48.3. The EAR module

Describe packaging with assembly descriptors

48.3.1. Assembly descriptor

Example 48.1. Sample of customized assembly for a standard EAR, configured for development

```
            <!-- This is a template assembly file that generates a mono Nuxeo-EP server.
  Configuration is designed for development usage -->

<assembly>
  <!-- RESOURCES -->
  <assemble>
    <set>resources</set>
    <outputFile>/</outputFile>
  </assemble>
  <zipEntrySet id="resources">
    <artifact>
      org.nuxeo.ecm.platform:nuxeo-platform-ear:${nuxeo.platform.version}:zip:resources-mono
    </artifact>
  </zipEntrySet>

  <!-- ARTIFACTS -->
  <assemble>
    <outputFile>/</outputFile>
    <set>root-artifacts</set>
  </assemble>
  <artifactSet id="root-artifacts">
    <import>**</import>
    <includeDependencies>false</includeDependencies>
    <includes>
      <artifact name="nuxeo-platform-webapp" />
      <artifact name="nuxeo-platform-webapp-core" />
    </includes>
  </artifactSet>

  <assemble>
    <outputFile>system</outputFile>
    <set>system</set>
  </assemble>
  <artifactSet id="system">
    <import>**</import>
    <includeDependencies>true</includeDependencies>
    <includes>
      <artifact group="org.nuxeo.*" category="runtime,jboss4"
        includeDependsOnCategory="false" />
      <artifact group="org.nuxeo.*" category="core,search,web" />
      <!-- add here nuxeo-addons -->
    </includes>
    <excludes>
      <artifact group="!org.nuxeo*" />
      <artifact group="org.nuxeo.common" />
      <artifact name="nuxeo-runtime-jboss-extensions" />

      <artifact name="nuxeo-platform-webapp" />
      <artifact name="nuxeo-platform-webapp-core" />

      <artifact name="nuxeo-platform-audit-facade" />
      <artifact name="nuxeo-platform-placeful-facade" />
      <artifact name="nuxeo-platform-search-compass-plugin" />
      <artifact name="nuxeo-platform-ear" />
      <artifact name="nuxeo-apt-extensions" />
      <artifact group="org.nuxeo.projects.template" />
    </excludes>
  </artifactSet>

  <assemble>
```

```
    <outputFile>system</outputFile>
    <unpack>true</unpack>
    <unpackInNewDirectory>true</unpackInNewDirectory>
    <set>nuxeo-platform-unpacked</set>
  </assemble>
  <artifactSet id="nuxeo-platform-unpacked">
    <import>**</import>
    <includes>
      <artifact name="nuxeo-platform-audit-facade" />
      <artifact name="nuxeo-platform-placeful-facade" />
      <artifact name="nuxeo-platform-search-compass-plugin" />
    </includes>
  </artifactSet>

  <!-- third party libraries embedded in the ear -->
  <assemble>
    <outputFile>lib</outputFile>
    <set>nuxeo-fixed-libs</set>
  </assemble>
  <artifactSet id="nuxeo-fixed-libs">
    <artifacts>
      <artifact group="org.freemarker" name="freemarker" version="2.3.11" />
      <artifact group="org.osgi" name="osgi-core" version="4.1" />
      <artifact group="commons-collections" name="commons-collections"
        version="3.1" />
      <artifact group="commons-io" name="commons-io" version="1.2" />
      <artifact group="commons-lang" name="commons-lang" version="2.2" />
      <artifact group="commons-fileupload" name="commons-fileupload"
        version="1.1.1" />
      <artifact group="cssparser" name="cssparser" version="0.9.4-fix" />
      <artifact group="net.sf.ehcache" name="ehcache" version="1.2.3" />
      <artifact group="net.sf.ezmorph" name="ezmorph" version="0.9" />
      <artifact group="org.hibernate" name="hibernate" version="3.2.0.ga" />
      <artifact group="org.hibernate" name="hibernate-annotations"
        version="3.2.0.ga" />
      <artifact group="org.hibernate" name="hibernate-entitymanager"
        version="3.2.0.ga" />
      <artifact group="jboss" name="jboss-cache-jdk50" version="1.4.0.SP1" />
      <artifact group="org.jboss.seam" name="jboss-seam" version="1.1.5.NX3" />
      <artifact group="jboss" name="jbpm" version="3.1.2" />
      <artifact group="jboss" name="jgroups" version="2.2.9" />
      <artifact group="net.sf.json-lib" name="json-lib" version="0.9" />
      <artifact group="org.apache.lucene" name="lucene-core" version="2.0.0" />
      <artifact group="net.sf.opencsv" name="opencsv" version="1.7" />
      <artifact group="org.slf4j" name="slf4j-api" version="1.3.0" />
      <artifact group="org.slf4j" name="slf4j-log4j12" version="1.3.0" />
      <artifact group="org.apache.myfaces.tomahawk" name="tomahawk"
        version="1.1.5" />
      <artifact group="org.apache.myfaces.tomahawk" name="tomahawk-sandbox"
        version="1.1.5" />
      <artifact group="org.apache.myfaces.trinidad" name="trinidad-api"
        version="1.0.1-incubating-NXEP51M2" />
      <artifact group="org.apache.myfaces.trinidad" name="trinidad-impl"
        version="1.0.1-incubating-NXEP51M2" />
      <artifact group="org.apache.directory.server"
        name="apacheds-protocol-shared" version="1.5.1" />
      <artifact group="org.apache.directory.shared" name="shared-ldap"
        version="0.9.7" />
      <artifact group="com.sun.facelets" name="jsf-facelets" version="1.1.11" />
      <artifact group="org.nuxeo.ecm.platform" name="nuxeo-jbossws-wrapper"
        version="4.0.5.GA" />
    </artifacts>
  </artifactSet>

  <!-- template project's artifacts -->
  <assemble>
    <outputFile>plugins</outputFile>
    <set>template-plugins</set>
  </assemble>
  <artifactSet id="template-plugins">
    <import>**</import>
    <includeDependencies>false</includeDependencies>
    <includes>
      <artifact group="org.nuxeo.projects.template" />
    </includes>
  </artifactSet>

  <!-- template project's resources -->
  <assemble>
    <outputFile>/</outputFile>
    <set>template-resources</set>
  </assemble>
  <fileSet id="template-resources">
    <directory>src/main/resources_template_common</directory>
    <excludes>
      <exclude>README.txt</exclude>
    </excludes>
  </fileSet>
  <assemble>
    <outputFile>/</outputFile>
    <set>template-resources-dev</set>
  </assemble>
  <fileSet id="template-resources-dev">
    <directory>src/main/resources_template_dev</directory>
```

```
    <excludes>
      <exclude>README.txt</exclude>
    </excludes>
  </fileSet>

</assembly>
```

Create an XML file in `template-ear/src/main/assemble/` that will describe the EAR to build. Here's the file content:

- assembly: the main tag. For now, you can only define one assembly per file (meaning one packaging definition by assembly descriptor file).
- assemble: part of an assembly. An assemble must be associated with a set (different sets are available: zipEntrySet, fileSet, artifactSet).
  - set: string. Id of associated set.
  - outputFile: output directory string. "/" represents the building EAR root.
  - unpack: true or false. Whether the set content should be unpack.
  - unpackInNewDirectory: true or false. if unpack is true, whether to create new directories or not in the output directory.
- zipEntrySet

  This set allow to retrieve artifacts from outside the project dependencies and unzip them.
  - id: required set id.
  - artifact: group:name:version:type:classifier of an artifact to retrieve from maven resolution, even if not present in the project's dependency tree. Will be unzipped if necessary.
  - profile: if specified, will only be treated if the given maven profile is active.
- artifactSet

  This set retrieves all artifacts from the project's dependencies tree, depending on conditions
  - id: required set id.
  - import: "**" is required for inheriting project's dependencies.
  - includeDependencies: true or false. If true, all dependencies of selected artifacts by this set will be added to the set.
  - excludeDependencies: true or false. If true, all dependencies of selected artifacts by this set will be removed from the set.
  - extends: deprecated. Was aimed to allow assemblies inheritance and sets overriding.
  - includes (or excludes): artifacts list to include (resp. exclude). Each artifact accept these parameters:
    - group: artifact's group id. Accept wildcards '*' and negative '!' to match multiple artifacts.
    - name: artifact's id. Accept wildcards '*' and negative '!' to match multiple artifacts.
    - type: artifact's type (pom, jar, ejb, ...).
    - version: artifact's version.

    - scope: artifact's scope (test, compile, provided, ...).
    - classifier: artifact's classifier.
    - file: artifact's filename. Accept wildcards '*' and negative '!' to match multiple artifacts.
    - category: artifact's category. Concept introduced by Nuxeo to "tag" an artifact with a "category". This is done in the artifact's Manifest as "Bundle-Category". Multiple categories, comma separated, is interpreted as an union of matching artifacts.
    - includeDependsOnCategory: true or false. In case of category specified, whether to include or not any artifact which has at least one dependency matching the specified category.
    - profile: if specified, the artifact will only be treated (included or excluded) if the given maven profile is active.
- fileSet

  This set is for manipulating local resources.

    - id: required set id.
    - directory: root directory of resources to copy. For now, only one directory per set is allowed
    - includes: files and directories list to include.
        - include: string. File, directory or pattern file/directory to include. "**" means multiple directories.
    - excludes: files and directories list to exclude. By default, these patterns are already excluded: "**/.svn" and "**/.hg".
        - exclude: string. File, directory or pattern file/directory to exclude. "**" means multiple directories.

48.3.2. Add some resources

With this assembly descriptor file, you can configure your local resources to be copied into the packaged EAR. The above template gives some usual resources folders.

- resources_template_common where you put your resources that should always be added or overwrite the default ones.
- resources_template_dev where you put your resources configured for development purpose (filesystem backend, ...).
- resources_template_ldap_pg is an example where you can put resources parameterized to configure Nuxeo with LDAP users and/or groups and a PostgreSQL backend.
- resources_template_stateful and resources_template_stateless are to be used for a bi-machine packaging, explained later in this chapter.

48.4. Improve usability

It is important to give simple and stable ways to use these assembly descriptors. Maven will allow to configure profiles and parameters. Ant will then give possibility to call multiple maven and shell commands, OS independent.

48.4.1. Thanks to Maven

The module-ear `pom.xml` file must have in its dependencies all wanted artifacts that will then be filtered by the assembly descriptor. Including nuxeo-platform-ear bring all its dependencies, there's no need to copy Nuxeo's artifacts list.

In order to call the assembly descriptor, it must contain the following plugin declaration. "template.ear.assembly" is used to parameterize the assembly descriptor name. Set a default value with a property.

```
  <properties>
    <!-- default assembly descriptor to use -->
    <template.ear.assembly>template-dev</template.ear.assembly>
  </properties>

  <build>
    <plugins>
      <plugin>
        <groupId>org.apache.maven.plugins</groupId>
        <artifactId>maven-nuxeo-plugin</artifactId>
        <extensions>true</extensions>
        <configuration>
          <runPreprocessor>false</runPreprocessor>
          <format>directory</format>
          <outputDirectory>../target</outputDirectory>
          <targetFile>${template.ear.assembly}</targetFile>
          <descriptor>
            ${basedir}/src/main/assemble/${template.ear.assembly}.xml
          </descriptor>
        </configuration>
        <executions>
          <execution>
            <id>assemble-ear</id>
            <phase>package</phase>
            <goals>
              <goal>assembly</goal>
            </goals>
          </execution>
        </executions>
      </plugin>
    </plugins>
  </build>
```

Maven parameters may be set in command-line. For example, in the above template, "mvn clean package" will package a Nuxeo EAR for developers and "mvn clean package -Dtemplate.ear.assembly=template-ldap-pg" will call template-ldap-pg.xml instead of template-dev.xml and so package a Nuxeo EAR configured for LDAP and PostgreSQL.

You can also use Maven profiles with "-P" argument.

48.4.2. Thanks to Ant

Using Ant, we will automatize lot of common tasks. Here's the output of command "ant -projecthelp" on template-project:

```
$ ant -projecthelp
Buildfile: build.xml

Main targets:

 copy             Replace ear and copy libs to jboss
 copy-2parts      Copy Nuxeo-EP in two parts
 copy-3parts      DEPRECATED: Copy Nuxeo-EP in three parts
 copy-ldap-pg     Copy template on LDAP and Postgres
 deploy           Build and copy to JBoss
 deploy-2parts    Build 2parts and copy to two patched JBoss
 deploy-ldap-pg   Build for LDAP and Postgres and copy to JBoss
 install          Run mvn install
 package          Package Nuxeo-EP
 package-2parts   Package Nuxeo-EP in two parts
 package-3parts   DEPRECATED: Package Nuxeo-EP in three parts
 package-ldap-pg  Package for LDAP and Postgres
 web              Copy web files to a live JBoss
Default target: usage
```

Example 48.2. build.xml sample

```
             <?xml version="1.0"?>
<project name="template" default="usage" basedir=".">

  <!-- Create a build.properties file from build.properties.sample
       if you wish to override the JBoss paths -->
  <property file="build.properties" />
  <property name="jboss.dir" value="/opt/jboss" />
  <property name="jboss.config" value="default" />
  <property name="mvn.opts" value="" />

  <property name="javac.debug" value="true" />
  <property name="javac.deprecation" value="false" />

  <!-- Boilerplate configuration -->
  <property name="build.dir" value="${basedir}/target" />
  <property name="deploy.dir" value="${jboss.dir}/server/${jboss.config}/deploy" />
  <property name="template.ear.root" value="template-ear" />
  <property name="nuxeo.ear" value="nuxeo.ear" />
  <property name="deploy.lib.dir" value="${jboss.dir}/server/${jboss.config}/lib" />
  <property name="build.lib.dir" value="lib" />


  <!-- these properties will need to be overridden at execution time -->
  <target name="set.jboss.home" unless="jboss.home">
    <property name="jboss.home" value="${jboss.dir}" />
  </target>
  <target name="set.assembly.name" unless="assembly.name">
    <property name="assembly.name" value="template-dev" />
  </target>
  <target name="setproperties" unless="jboss.server" depends="set.jboss.home, set.assembly.name">
    <property name="assembly.ear" value="${assembly.name}" />
    <property name="jboss.server" value="${jboss.home}/server/${jboss.config}" />
    <property name="jboss.deploy" value="${jboss.server}/deploy" />
    <property name="jboss.lib" value="${jboss.server}/lib" />
    <property name="jboss.nuxeo.ear" value="${jboss.deploy}/nuxeo.ear" />
    <property name="template.ear.build" value="${template.ear.root}/target/${assembly.ear}" />
  </target>

  <target name="usage">
    <echo message="usage: ant [package|deploy|web|package-2parts|deploy-2parts|package-ldap-pg|deploy-ldap-p]" /
    <echo message="ant package         => Package template to ${build.dir}/${nuxeo.ear}" />
    <echo message="ant deploy          => Package template and copy to ${deploy.dir}" />
    <echo message="ant deploy-Nparts   => Package template and copy to ${deploy.dir} relative to the ${jbossN.di
    <echo message="ant web             => Copy web files to a live JBoss (${deploy.dir})" />
  </target>

  <condition property="osfamily-unix">
    <os family="unix" />
  </condition>
  <condition property="osfamily-windows">
    <os family="windows" />
  </condition>

  <target name="deploy" depends="install,package,copy,copy-lib" description="Build and copy to JBoss">
    <echo message="Deployed ${ant.project.name} to ${deploy.dir}" />
  </target>

  <target name="deploy-ldap-pg" depends="install,package-ldap-pg,copy-ldap-pg" description="Build for LDAP and P
  </target>
  <target name="deploy-2parts" depends="install,package-2parts,copy-2parts" description="Build 2parts and copy t
  </target>

  <target name="package-ldap-pg" description="Package for LDAP and Postgres">
    <antcall target="package">
      <param name="assembly.name" value="template-ldap-pg" />
    </antcall>
  </target>
  <target name="package-2parts" description="Package Nuxeo-EP in two parts">
    <antcall target="package">
      <param name="jboss.home" value="${jboss1.dir}" />
      <param name="assembly.name" value="template-stateful" />
    </antcall>
    <antcall target="package">
      <param name="jboss.home" value="${jboss2.dir}" />
      <param name="assembly.name" value="template-stateless" />
    </antcall>
  </target>
  <target name="package-3parts" description="DEPRECATED: Package Nuxeo-EP in three parts">
    <antcall target="package">
      <param name="assembly.name" value="nuxeo-core" />
    </antcall>
    <antcall target="package">
      <param name="assembly.name" value="nuxeo-indexing" />
    </antcall>
    <antcall target="package">
      <param name="assembly.name" value="nuxeo-webplatform" />
    </antcall>
  </target>
  <target name="package" depends="setproperties,package-unix,package-windows" description="Package Nuxeo-EP" />
  <target name="package-unix" if="osfamily-unix">
    <echo message="assembly NAME ${assembly.name}" />
    <exec executable="mvn" failonerror="true">
      <arg value="clean" />
      <arg value="package" />
```

```
    <arg value="-Dmaven.test.skip=true" />
    <arg value="-f" />
    <arg value="${template.ear.root}/pom.xml" />
    <arg value="-Dtemplate.ear.assembly=${assembly.name}" />
    <arg value="${mvn.opts}" />
  </exec>
  <echo message="Packaged ${ant.project.name} into ${template.ear.build}" />
</target>
<target name="package-windows" if="osfamily-windows">
  <exec executable="cmd" failonerror="true">
    <arg value="/c" />
    <arg value="mvn.bat" />
    <arg value="clean" />
    <arg value="package" />
    <arg value="-Dmaven.test.skip=true" />
    <arg value="-f" />
    <arg value="${template.ear.root}/pom.xml" />
    <arg value="-Dtemplate.ear.assembly=${assembly.name}" />
    <arg value="${mvn.opts}" />
  </exec>
</target>
<target name="install" depends="install-unix,install-windows" description="Run mvn install" />
<target name="install-unix" if="osfamily-unix">
  <exec executable="mvn" failonerror="true">
    <arg value="clean" />
    <arg value="install" />
    <arg value="-Dmaven.test.skip=true" />
    <arg value="${mvn.opts}" />
  </exec>
</target>
<target name="install-windows" if="osfamily-windows">
  <exec executable="cmd" failonerror="true">
    <arg value="/c" />
    <arg value="mvn.bat" />
    <arg value="clean" />
    <arg value="install" />
    <arg value="-Dmaven.test.skip=true" />
    <arg value="${mvn.opts}" />
  </exec>
</target>

<target name="web" description="Copy web files to a live JBoss">
  <copy todir="${deploy.dir}/nuxeo.ear/nuxeo.war">
    <fileset dir="${basedir}/template-platform/src/main/resources/nuxeo.war/" />
  </copy>
</target>

<target name="rename">
  <delete dir="${build.dir}/${nuxeo.ear}" failonerror="false" />
  <move file="${build.dir}/template.ear" tofile="${build.dir}/${nuxeo.ear}" />
</target>

<target name="copy-lib">
  <delete failonerror="false">
    <fileset dir="${deploy.lib.dir}/">
      <include name="nuxeo*.jar" />
    </fileset>
  </delete>
  <copy todir="${deploy.lib.dir}/" failonerror="true">
    <fileset dir="${build.dir}/${build.lib.dir}/" />
  </copy>
</target>

<target name="copy-ldap-pg" description="Copy template on LDAP and Postgres">
  <antcall target="copy">
    <param name="assembly.name" value="template-ldap-pg" />
  </antcall>
</target>
<target name="copy-2parts" description="Copy Nuxeo-EP in two parts">
  <antcall target="copy">
    <param name="jboss.home" value="${jboss1.dir}" />
    <param name="assembly.name" value="template-platform-stateful" />
  </antcall>
  <antcall target="copy">
    <param name="jboss.home" value="${jboss2.dir}" />
    <param name="assembly.name" value="template-web-stateless" />
  </antcall>
</target>
<target name="copy-3parts" description="DEPRECATED: Copy Nuxeo-EP in three parts">
  <antcall target="copy">
    <param name="jboss.home" value="${jboss1.dir}" />
    <param name="assembly.name" value="nuxeo-core" />
  </antcall>
  <antcall target="copy">
    <param name="jboss.home" value="${jboss2.dir}" />
    <param name="assembly.name" value="nuxeo-indexing" />
  </antcall>
  <antcall target="copy">
    <param name="jboss.home" value="${jboss3.dir}" />
    <param name="assembly.name" value="nuxeo-webplatform" />
  </antcall>
</target>
<target name="copy" depends="delete-ear,copy-ear,copy-lib" description="Replace ear and copy libs to jboss" />
<target name="delete-ear" depends="setproperties">
```

```
    <delete dir="${jboss.nuxeo.ear}" failonerror="false" />
  </target>
  <target name="copy-ear" depends="setproperties">
    <mkdir dir="${jboss.nuxeo.ear}" />
    <copy todir="${jboss.nuxeo.ear}">
      <fileset dir="${template.ear.build}" />
    </copy>
  </target>

  <target name="copy-jars">
    <copy todir="${deploy.dir}/${nuxeo.ear}/plugins" overwrite="true" flatten="true">
      <fileset dir="${basedir}">
        <include name="*/target/*.jar" />
        <exclude name="*/target/*-sources.jar" />
      </fileset>
    </copy>
  </target>

</project>
```

48.5. Recommended multi-machine packagings

48.5.1. Bi-machine: stateful/stateless

48.5.1.1. Introduction

Stateful/Stateless packaging gives a good solution to a lot of production requirements. Principle is to split Nuxeo in two parts: one that contains the core and every services that persist data, and another that provides web services. Stateless part may be duplicated.

48.5.1.2. Packaging

Bi-machine packaging is done by two assembly descriptors, each one building an EAR to deploy on a separate server.

To build your own bi-machine packaging, copy the template-project structure. Copy and adapt Ant, Maven and descriptor files. Then, configure artifactSet and fileSet that will deploy your project's artifacts into the plugin directory and the wanted resources to replace those deployed by default.

Example 48.3. Sample of customized assembly for a stateful EAR

```
             <!-- This is a template assembly file that generates a "stateful" part of Nuxeo-EP server -->
<assembly>

  <!-- RESOURCES -->
  <assemble>
    <set>resources</set>
    <outputFile>/</outputFile>
  </assemble>
  <zipEntrySet id="resources">
    <artifact>
      org.nuxeo.ecm.platform:nuxeo-platform-ear:${nuxeo.platform.version}:zip:resources-platform-stateful
    </artifact>
  </zipEntrySet>

  <!-- ARTIFACTS -->
  <assemble>
    <outputFile>system</outputFile>
    <set>system</set>
  </assemble>
  <artifactSet id="system">
    <import>**</import>
    <includeDependencies>true</includeDependencies>
    <includes>
      <artifact group="org.nuxeo.*" category="runtime,jboss4"
        includeDependsOnCategory="false" />
      <artifact group="org.nuxeo.*" category="core,stateful" />
      <!-- add here core and stateful nuxeo-addons -->
```

```
    </includes>
    <excludes>
      <artifact group="!org.nuxeo*" />
      <artifact group="org.nuxeo.common" />
      <artifact name="nuxeo-runtime-jboss-extensions" />

      <artifact name="nuxeo-platform-webapp" />
      <artifact name="nuxeo-platform-webapp-core" />

      <artifact name="nuxeo-platform-audit-facade" />
      <artifact name="nuxeo-platform-placeful-facade" />
      <artifact name="nuxeo-platform-search-compass-plugin" />
      <artifact name="nuxeo-platform-ear" />
      <artifact group="org.nuxeo.projects.template" />
    </excludes>
  </artifactSet>

  <assemble>
    <outputFile>system</outputFile>
    <unpack>true</unpack>
    <unpackInNewDirectory>true</unpackInNewDirectory>
    <set>nuxeo-platform-unpacked</set>
  </assemble>
  <artifactSet id="nuxeo-platform-unpacked">
    <import>**</import>
    <includes>
      <artifact name="nuxeo-platform-audit-facade" />
      <artifact name="nuxeo-platform-placeful-facade" />
      <artifact name="nuxeo-platform-search-compass-plugin" />
    </includes>
  </artifactSet>

  <!-- third party libraries embedded in the ear -->
  <assemble>
    <outputFile>lib</outputFile>
    <set>nuxeo-fixed-libs</set>
  </assemble>
  <artifactSet id="nuxeo-fixed-libs">
    <artifacts>
      <artifact group="org.freemarker" name="freemarker" version="2.3.11" />
      <artifact group="org.osgi" name="osgi-core" version="4.1" />
      <artifact group="commons-collections" name="commons-collections"
        version="3.1" />
      <artifact group="commons-io" name="commons-io" version="1.2" />
      <artifact group="commons-lang" name="commons-lang" version="2.2" />
      <artifact group="commons-fileupload" name="commons-fileupload"
        version="1.1.1" />
      <artifact group="cssparser" name="cssparser" version="0.9.4-fix" />
      <artifact group="net.sf.ehcache" name="ehcache" version="1.2.3" />
      <artifact group="net.sf.ezmorph" name="ezmorph" version="0.9" />
      <artifact group="org.hibernate" name="hibernate" version="3.2.0.ga" />
      <artifact group="org.hibernate" name="hibernate-annotations"
        version="3.2.0.ga" />
      <artifact group="org.hibernate" name="hibernate-entitymanager"
        version="3.2.0.ga" />
      <artifact group="jboss" name="jboss-cache-jdk50" version="1.4.0.SP1" />
      <artifact group="org.jboss.seam" name="jboss-seam" version="1.1.5.NX3" />
      <artifact group="jboss" name="jbpm" version="3.1.2" />
      <artifact group="jboss" name="jgroups" version="2.2.9" />
      <artifact group="net.sf.json-lib" name="json-lib" version="0.9" />
      <artifact group="org.apache.lucene" name="lucene-core" version="2.0.0" />
      <artifact group="net.sf.opencsv" name="opencsv" version="1.7" />
      <artifact group="org.slf4j" name="slf4j-api" version="1.3.0" />
      <artifact group="org.slf4j" name="slf4j-log4j12" version="1.3.0" />
      <artifact group="org.apache.myfaces.tomahawk" name="tomahawk"
        version="1.1.5" />
      <artifact group="org.apache.myfaces.tomahawk" name="tomahawk-sandbox"
        version="1.1.5" />
      <artifact group="org.apache.myfaces.trinidad" name="trinidad-api"
        version="1.0.1-incubating-NXEP51M2" />
      <artifact group="org.apache.myfaces.trinidad" name="trinidad-impl"
        version="1.0.1-incubating-NXEP51M2" />
      <artifact group="org.apache.directory.server"
        name="apacheds-protocol-shared" version="1.5.1" />
      <artifact group="org.apache.directory.shared" name="shared-ldap"
        version="0.9.7" />
      <artifact group="com.sun.facelets" name="jsf-facelets" version="1.1.11" />
      <artifact group="org.nuxeo.ecm.platform" name="nuxeo-jbossws-wrapper"
        version="4.0.5.GA" />
    </artifacts>
  </artifactSet>

  <!-- template project's artifacts -->
  <assemble>
    <outputFile>plugins</outputFile>
    <set>template-plugins</set>
  </assemble>
  <artifactSet id="template-plugins">
    <import>**</import>
    <includeDependencies>true</includeDependencies>
    <includes>
        <artifact name="template-core" />
        <artifact name="template-stateful-services" />
    </includes>
```

```
    <excludes>
      <artifact group="!org.nuxeo.projects.template" />
    </excludes>
  </artifactSet>

  <!-- template project's resources -->
  <assemble>
    <outputFile>/</outputFile>
    <set>template-resources</set>
  </assemble>
  <fileSet id="template-resources">
    <directory>src/main/resources_template_common</directory>
    <excludes>
      <exclude>README.txt</exclude>
    </excludes>
  </fileSet>
  <assemble>
    <outputFile>/</outputFile>
    <set>template-resources-stateful</set>
  </assemble>
  <fileSet id="template-resources-stateful">
    <directory>src/main/resources_template_stateful</directory>
  </fileSet>

</assembly>
```

Example 48.4. Sample of customized assembly for a stateless EAR

```
              <!-- This is a template assembly file that generates a "stateless" part of Nuxeo-EP server -->

<assembly>

  <!-- RESOURCES -->
  <assemble>
    <set>resources</set>
    <outputFile>/</outputFile>
  </assemble>
  <zipEntrySet id="resources">
    <artifact>
      org.nuxeo.ecm.platform:nuxeo-platform-ear:${nuxeo.platform.version}:zip:resources-web-stateless
    </artifact>
  </zipEntrySet>

  <!-- ARTIFACTS -->
  <assemble>
    <outputFile>/</outputFile>
    <set>root-artifacts</set>
  </assemble>
  <artifactSet id="root-artifacts">
    <import>**</import>
    <includeDependencies>false</includeDependencies>
    <includes>
      <artifact name="nuxeo-platform-webapp" />
      <artifact name="nuxeo-platform-webapp-core" />
    </includes>
  </artifactSet>

  <assemble>
    <outputFile>system</outputFile>
    <set>system</set>
  </assemble>
  <artifactSet id="system">
    <import>**</import>
    <includeDependencies>true</includeDependencies>
    <includes>
      <artifact group="org.nuxeo.*" category="runtime,jboss4"
        includeDependsOnCategory="false" />
      <artifact group="org.nuxeo.*" category="stateless" />
      <!-- add here web and stateless nuxeo-addons -->
    </includes>
    <excludes>
      <artifact group="!org.nuxeo*" />
      <artifact group="org.nuxeo.common" />
      <artifact name="nuxeo-runtime-jboss-extensions" />

      <artifact name="nuxeo-platform-webapp" />
      <artifact name="nuxeo-platform-webapp-core" />
      <artifact name="nuxeo-platform-ear" />
      <artifact name="nuxeo-apt-extensions" />
      <artifact group="org.nuxeo.projects.template" />
    </excludes>
  </artifactSet>

  <!-- third party libraries embedded in the ear -->
  <assemble>
```

```
    <outputFile>lib</outputFile>
    <set>nuxeo-fixed-libs</set>
  </assemble>
  <artifactSet id="nuxeo-fixed-libs">
    <artifacts>
      <artifact group="org.freemarker" name="freemarker" version="2.3.11" />
      <artifact group="org.osgi" name="osgi-core" version="4.1" />
      <artifact group="commons-collections" name="commons-collections"
        version="3.1" />
      <artifact group="commons-io" name="commons-io" version="1.2" />
      <artifact group="commons-lang" name="commons-lang" version="2.2" />
      <artifact group="commons-fileupload" name="commons-fileupload"
        version="1.1.1" />
      <artifact group="cssparser" name="cssparser" version="0.9.4-fix" />
      <artifact group="net.sf.ehcache" name="ehcache" version="1.2.3" />
      <artifact group="net.sf.ezmorph" name="ezmorph" version="0.9" />
      <artifact group="org.hibernate" name="hibernate" version="3.2.0.ga" />
      <artifact group="org.hibernate" name="hibernate-annotations"
        version="3.2.0.ga" />
      <artifact group="org.hibernate" name="hibernate-entitymanager"
        version="3.2.0.ga" />
      <artifact group="jboss" name="jboss-cache-jdk50" version="1.4.0.SP1" />
      <artifact group="org.jboss.seam" name="jboss-seam" version="1.1.5.NX3" />
      <artifact group="jboss" name="jbpm" version="3.1.2" />
      <artifact group="jboss" name="jgroups" version="2.2.9" />
      <artifact group="net.sf.json-lib" name="json-lib" version="0.9" />
      <artifact group="org.apache.lucene" name="lucene-core" version="2.0.0" />
      <artifact group="net.sf.opencsv" name="opencsv" version="1.7" />
      <artifact group="org.slf4j" name="slf4j-api" version="1.3.0" />
      <artifact group="org.slf4j" name="slf4j-log4j12" version="1.3.0" />
      <artifact group="org.apache.myfaces.tomahawk" name="tomahawk"
        version="1.1.5" />
      <artifact group="org.apache.myfaces.tomahawk" name="tomahawk-sandbox"
        version="1.1.5" />
      <artifact group="org.apache.myfaces.trinidad" name="trinidad-api"
        version="1.0.1-incubating-NXEP51M2" />
      <artifact group="org.apache.myfaces.trinidad" name="trinidad-impl"
        version="1.0.1-incubating-NXEP51M2" />
      <artifact group="org.apache.directory.server"
        name="apacheds-protocol-shared" version="1.5.1" />
      <artifact group="org.apache.directory.shared" name="shared-ldap"
        version="0.9.7" />
      <artifact group="com.sun.facelets" name="jsf-facelets" version="1.1.11" />
      <artifact group="org.nuxeo.ecm.platform" name="nuxeo-jbossws-wrapper"
        version="4.0.5.GA" />
    </artifacts>
  </artifactSet>

  <!-- template project's artifacts -->
  <assemble>
    <outputFile>plugins</outputFile>
    <set>template-plugins</set>
  </assemble>
  <artifactSet id="template-plugins">
    <import>**</import>
    <includeDependencies>true</includeDependencies>
    <includes>
        <artifact name="template-web" />
        <artifact name="template-stateless-services" />
    </includes>
    <excludes>
      <artifact group="!org.nuxeo.projects.template" />
    </excludes>
  </artifactSet>

  <!-- template project's resources -->
  <assemble>
    <outputFile>/</outputFile>
    <set>template-resources</set>
  </assemble>
  <fileSet id="template-resources">
    <directory>src/main/resources_template_common</directory>
    <excludes>
      <exclude>README.txt</exclude>
    </excludes>
  </fileSet>
  <assemble>
    <outputFile>/</outputFile>
    <set>template-resources-stateless</set>
  </assemble>
  <fileSet id="template-resources-stateless">
    <directory>src/main/resources_template_stateless</directory>
  </fileSet>

</assembly>
```

48.5.1.3. Network configuration

Stateful is identified in configuration files as **nxplatformserver** and **nxjmsserver**. Stateless is known as **nxwebserver**. You can set this in your hosts files.

If you deploy multiple stateless servers, the unique (for now it cannot be duplicated) stateful server only need to know one stateless server. Stateless servers do not need to communicate between each other.

Note

There is misconfiguration in Nuxeo EP 5.1.6 and 5.2.M3 (http://jira.nuxeo.org/browse/NXP-2838). Replace in `nuxeo-web-stateless.ear/datasources/core-events-ds.xml`

```
java.naming.provider.url=${jboss.bind.address}:1099
```

with

```
java.naming.provider.url=nxjmsserver:1099
```

Chapter 49. Release Management

49.1. Introduction

Nuxeo EP and related components can use the maven release plugin to prepare and perform releases of Nuxeo components.

The release of a project usually consists in those operations:

- perform a clean checkout and launch the tests
- change version numbers
- tag sources in the source control system (ex: SVN)
- update version number of the trunk to start the new development cycle
- publish software and related information

The maven release plugin streamlines and automates the release process of software components.

To get more background about maven release, you can read following pages:

- mini-guide on Maven Release
- maven release main site

49.2. Let's release!

When your software is ready to be released, you want to launch the release process. :-)

This process has to be as easy as possible. You will be able to find bellow the basic process.

49.2.1. Remove all dependencies on SNAPSHOT versions

To enforce that the build is reproducible, maven refuses to create a new release if there are dependencies on SNAPSHOT versions.

The first step is then to remove all dependencies on SNAPSHOT versions. Of course, as you may not be able to release all your dependencies as long as your software, you can create specific versions of your dependencies for the release.

Here is the fastest way:

```
$ mvn deploy:deploy-file -DrepositoryId=central \ ❶
-Durl=scpexe://amour.nuxeo.org/home/nexus/repositories/vendor-release \ ❷
-DgroupId=org.apache.myfaces.trinidad \ ❸
-DartifactId=trinidad-impl \ ❹
-Dversion=1.0.1-incubating-NXEP51M2 \ ❺
-Dpackaging=jar \ ❻
-Dfile=trinidad-impl-1.0.1-incubating-NXEP51M2.jar ❼
```

❶❷ Informations of the repository you are using. For Nuxeo's software, use:

- `repositoryId`: `central`
- `url`: `scpexe://amour.nuxeo.org/home/nexus/repositories/vendor-release>`

Of course, you will need `write` permission on the repository. You can request it from your repository's manager.

❸❹❺❻
- `groupId`: the group id for the artifact, you should get it for the pom of this artifact
- `artifactId`: idem
- `packaging`: usually `jar` (or `pom` if you are only deploying a pom file)

❼ Point to the file (usually jar or pom) you want to upload.

49.2.2. Checkout the code and clean your repository

First, you need to get a fresh copy of the components you are about to release.

```
$ mvn scm:checkout svn:http://server/yoursoftware/trunk
```

To be sure that everything is clean and that you will test in the same environment as new users, remove your repository. On windows, remove your repository using your preferred tool. On Linux/Unix/MacOS, you can perform:

```
$ rm ~/.m2/repository
```

to remove your entire maven repository.

49.2.3. Test the release

Let's try to build and test the software before launching a dry run of the release process.

```
mvn install
```

This should download all required dependencies (since you've clean your maven repository in the previous paragraph) and build the software.

We now can test the release preparation process.

```
mvn release:prepare -DdryRun=true
```

Answer, now to the questions maven will ask about version numbers and tag name. This will simulate a release preparation, especially the svn tags.

If everything went well until that point, we can actually perform the release...

49.2.4. Perform the release

For each component you want to release, do:

```
mvn release:prepare release:perform
```

This will tag your sources in the SCM, build the software, build the site upload your artifact to the defined repository and deploy your site.

In Nuxeo's case, you need to release first the global master POM before releasing other component, if the master POM is using a SNAPSHOT version. To achieve this, go to a checkout of the master POM and do a non-recursive release using:

```
mvn -N release:prepare release:perform
```

Then you can set master POM's new version number into each component you want to release and start the release process.

49.2.5. You're done

Now you're done. You have released your software using maven release and you can happily use your new component in your projects or deliver your customer! :-)

Part VII. Add-ons

Chapter 50. Add-ons

50.1. Introduction

Add-ons are modules you can add to a Nuxeo instance and that provide additional functionalities. To install an add-ons, download the modules from http://download.nuxeo.org/ and copy them in the `/plugins` directory of your Nuxeo instance. For additional set-up, refer to the specific add-on's documentation.

Chapter 51. Nuxeo Annotation Service

51.1. Introduction

51.1.1. W3C Annotea

Nuxeo Annotation service is based on Annotea W3C specification. It is compliant with the specification published at Annotea Protocol .

However the annotea specification is quite old and does not handle all use cases for annotating documents in the context of Nuxeo ECM. Therefore, we propose to provide some extensions to annotea. Using the repository plugin, it provides a plugable system to be able to keep track of the relationship between an annotated URL and the underlying nuxeo document.

Inside Nuxeo EP we need to be able to annotate the HTML preview of a document. This preview may include images. This images are annotable as the rest of the document, but it requires an extension to the Xpointer specification since user may want to annotate just a portion (zone) of the image.

51.1.2. Logical architecture overview

This section provides an overview of the logical architecture of the Nuxeo Annotation service.

- Annotation Service API `nuxeo-platform-annotations-api`

  Provides Java API and all the needed Java artifacts needed to call the Annotation service remotely.

- Annotation Service Core `nuxeo-platform-annotations-core`

  Provides the Nuxeo Service that exposes the required annotation Java Interface. This service is implemented as a Nuxeo Runtime component and will provide the needed extension points .

- Annotation Service Facade `nuxeo-platform-annotations-facade`

  Provides the service EJB3 facade for remote (RMI) access.

- Annotation Service Http gateway `nuxeo-platform-annotations-http`

  The HTTP Gateway implements the Annotea HTTP protocol. It is implemented as a servlet that enable access to the annotation service via http GET/POST requests as defined in the W3C Annotea specification.

- Annotation Service Web Interface `nuxeo-platform-annotations-web`

  The web interface provides ways to annotate document.

51.1.3. XPointer integration and extension.

Annotea uses XPointer to annotates part of a web pages. We added a image-range function to XPointer to be able to annotates part of an image. A XPointer using image-range looks like: "http://example.com/foo.html#xpointer(image-range(/html[1]/body[0]/img[0],[79,133],[123,159]))". The form is "image-range(location-set , [tlX,tlY], [brX,brY])", with the location-set defined in the XPointer documentation, pointing to an image. tlX, tlY, brX and brY are integer for the top left X, top left Y, bottom right X and bottom right Y coordinates of a square within the annotated original image (as stored on the server). That is, if the size of the image is dynamically modified using Javascript, an image-range should always point to the same area of the image.

51.2. Annotation Service Core

51.2.1. Overview

This is the main component of NXAS, the one that contains all the logic for managing RDF based annotations. The Core Service is not document aware and manages only URIs.

URI, URL and URN

A URL is "http://foo.com/index.html", a URN doesn't use a protocol: "nuxeo:default:12345678". A URI is a URL or a URN.

This component is also responsible for exposing all the needed extension points that will be used for configuration and integration.

51.2.2. Implementation

The service is implemented as a Runtime service on top of a Nuxeo Runtime component. The runtime component provides the extension point mechanisms. The API provided by the service will target managing annotations on URLs.

As any Nuxeo Service, the Annotation Service is accessible via the Runtime lookup method :

```
Framework.getLocalService(AnnotationService.class)
```

Implementation Note

The core function of the annotation service is to add/remove/query annotations from the RDF Graph. It also has to provide security and add configurability. The implementation separates those 2 functions in 2 classes: The AnnotationServiceProxy class implements all the configuration, event management and security function and let the AnnotationServiceImpl provides the core RDF service.

51.2.3. Storage

The Annotation service stores the annotations as a RDF graph. Nuxeo Relation Service is responsible for storing and managing the RDF data. The service uses a graph named "annotations". The default contribution is:

```
<component name="org.nuxeo.ecm.annotations.graph">
  <require>org.nuxeo.ecm.platform.relations.jena</require>
  <extension  target="org.nuxeo.ecm.platform.relations.services.RelationService"
      point="graphs">
    <graph name="annotations" type="jena">
      <option name="backend">sql</option>
      <option name="databaseType">${org.nuxeo.ecm.sql.jena.databaseType}</option>
      <option name="databaseTransactionEnabled">
        ${org.nuxeo.ecm.sql.jena.databaseTransactionEnabled}
      </option>
      <option name="datasource">java:/nxrelations-default-jena</option>
      <namespaces>
        <namespace name="rdf">http://www.w3.org/1999/02/22-rdf-syntax-ns#</namespace>
        <namespace name="dcterms">http://purl.org/dc/terms/</namespace>
        <namespace name="nuxeo">http://www.nuxeo.org/document/uid/</namespace>
      </namespaces>
    </graph>
  </extension>
</component>
```

Refer to the Relation Service documentation (Chapter 22, *Relations*) for more information on configuration options.

51.2.4. uriResolver

The uriResolver extension point allows to contribute a class that is responsible for resolving uri. It maps URI sent by the client to URI stored in the graph. It also allows to translate a URI in a list of URIs when querying the graph.

This allows to consider annotations on different URLs to be considered as annotations on the same document, therefore being able to see those annotations on both URLs.

To contribute a uriResolver contribute to the extension point a class that implements the UriResolver interface. The default contribution is:

```
<extension
    target="org.nuxeo.ecm.platform.annotations.services.AnnotationsService"
    point="uriResolver">
  <urlResolver
     class="org.nuxeo.ecm.platform.annotations.service.DefaultUriResolver"/>
</extension>
```

The DefaultUriResolver maps URI without changing them, unless it includes "nuxeo/Annotations" in its URL. In such a case, it transform the URL in a URN of the form "urn:annotation:***".

51.2.5. urlPatternFilter

The urlPatternFilter extension point allows to contribute regular expression pattern to the list of allowed URL pattern or disallowed URL pattern. When a request is made to get/set annotations on an URL, the server check the list.

The check is done before any transformation on the URL. Because the check on URL is done very early in the processing of a request, to increase performance, it is recommended to use it whenever possible.

The filter follows Apache's mod_access convention to filter URL. The default contribution is:

```
<extension
    target="org.nuxeo.ecm.platform.annotations.services.AnnotationsService"
    point="urlPatternFilter">
  <urlPatternFilter order="Deny,Allow">
    <deny>.*</deny>
    <allow>.*</allow>
  </urlPatternFilter>
</extension>
```

51.2.6. metadata

The metadata extension point allows to add metadata to an annotations when it is created. The contribution needs to implement the MetadataMapper interface. It is passed the annotation to which it can add the metadata. The default contribution:

```
<extension
    target="org.nuxeo.ecm.platform.annotations.services.AnnotationsService"
    point="metadataMapper">
  <metadataMapper
     class="org.nuxeo.ecm.platform.annotations.service.DefaultMetadataMapper">
</extension>
```

This contribution adds the date and creator to the annotation. Note that you can add metadata and modify the annotation before sending it to the service. You could choose to add the creation time using the local time of the client and add it to the annotation without using the metadatMapper.

51.2.7. permissionManager and permissionMapper

The permissionManager extension point allows to contribute a class that checks before CRUD operations. The permissionMapper will maps permission name with operation.

The permissionMapper extension point allows to contribute the name of the permission to check before a CRUD operation. The default contribution is:

```
<extension
    target="org.nuxeo.ecm.platform.annotations.services.AnnotationsService"
    point="permissionMapper">
  <permissionMapper>
    <createAnnotation>createAnnotation</createAnnotation>
    <readAnnotation>readAnnotation</readAnnotation>
    <updateAnnotation>updateAnnotation</updateAnnotation>
    <deleteAnnotation>deleteAnnotation</deleteAnnotation>
  </permissionMapper>
</extension>
```

The permissionManager extension point allows to contribute a class to check before CRUD operations. The default permission manager allows all operations to everybody:

```
<extension
  target="org.nuxeo.ecm.platform.annotations.services.AnnotationsService"
  point="permissionManager">
  <permissionManager
    class="org.nuxeo.ecm.platform.annotations.service.DefaultPermissionManager"/>
</extension>
```

51.2.8. annotabilityManager

The annotabilityManager extension point allows to contribute a class to fine grain which documents can be annotated. An implementation needs to implement the interface `AnnotabilityManager` . The default annotability manager is:

```
<extension
    target="org.nuxeo.ecm.platform.annotations.services.AnnotationsService"
    point="annotabilityManager">
  <annotabilityManager
      class="org.nuxeo.ecm.platform.annotations.service.DefaultAnnotabilityManager"/>
</extension>
```

It allows annotations on all URL.

51.2.9. UID Generation

When an annotation is created it is given a UID. This extension point allows to contribute a UID generator. The default contribution is:

```
<extension
  target="org.nuxeo.ecm.platform.annotations.services.AnnotationsService"
  point="annotationIDGenerator">
  <IDGenerator class="org.nuxeo.ecm.platform.annotations.service.DefaultIDGenerator"/>
</extension>
```

This default contribution create a UID using java `java.util.UUID.randomUUID()` .

51.2.10. Event management

The eventManager extension point allows to contribute a listener class to handle annotation related events. The listener class has to implement the interface. `EventListener` . There is one method for each CRUD operation, before and after the operation. The calls are made synchronously. There is no default contribution to this extension point.

51.3. NXAS Repository Plugin

51.3.1. Overview

NXAS core offers an implementation of an Annotea Server. It is "pure" RDF and relies on URI for resource identification. However, Nuxeo deals with versionned document. The repository plugin allows to contribute to the core module using Nuxeo's object of DocumentModel, ACP and Version.

The repository plugin contribute to Core NXAS, replacing most of the default contribution. You can then contribute to the repository plugin to tune security, annotability and more.

Querying using a DocumentModel

The Annotation Service uses URI exclusively, if you need to query using a DocumentModel, you need to obtain the URN for this document:

```
URNDocumentViewTranslator translator = new URNDocumentViewTranslator();
String documentId = documentModel.getId();
String serverName = coreSession.getRepositoryName();
URI uri = translator.getNuxeoUrn(serverName, documentId);
```

You can then use the uri for the query method.

51.3.2. Default contribution

This section describes the default contribution made to the annotation service. You should not change them if you use the repository plugin. To modify the repository plugin behavior, use its extension point (Section 51.3.3.1, “ documentAnnotability ”).

51.3.2.1. permissionMapper

The permission are mapped to the document that is being annotated:

```
<extension
  target="org.nuxeo.ecm.platform.annotations.services.AnnotationsService"
  point="permissionMapper">
  <permissionMapper>
    <createAnnotation>updateDocument</createAnnotation>
    <readAnnotation>viewDocument</readAnnotation>
    <updateAnnotation>updateDocument</updateAnnotation>
    <deleteAnnotation>deleteDocument</deleteAnnotation>
  </permissionMapper>
</extension>
```

51.3.2.2. uriResolver

The defaultNuxeoUriResolver transform the URI if it points to a Nuxeo document. The URI becomes a URN using the docId and repository name.

51.3.2.3. metadataMapper

In addition to the default mapper, the DefaultNuxeoMetadataMapper add the company of the user.

51.3.2.4. permissionMapper

The name of permission are the same as the document permission.

51.3.2.5. annotabilityManager

The contribution to the annotability manager will call the documentAnnotabilityManager.

51.3.2.6. eventListener

The contribution to the eventListener will call the corresponding AnnotatedDocumentEventListener.

51.3.3. Extension Point

51.3.3.1. documentAnnotability

The documentAnnotability extension point allows to choose which document are annotable. To contribute to this point, create a class that implement `DocumentAnnotability` . The default implementation is (`DefaultDocumentAnnotability`), allows annotations on all documents.

51.3.4. documentEventListener

The documentEventListener allows to contribute class that would be notify when an annotation is create, updated, read and deleted. Notification comes synchronously before and after the event. It is passed as parameter the annotation and the DocumentLocation of the the document being annoted.

A default contribution notify the JMS bus.

51.3.5. securityManager

The security manager extension point allows to contribute a class that implements the SecurityManager interface. It fine grains security using NuxeoPrincipal, the permission checked and the DocumentMode.

51.3.6. jcrLifecycleEventId and graphManagerEventListener

The jcrLifecycleEventId and graphManagerEventListener allows to modify the RDF graph when an event happens in the repository. The jcrLifecycleEventId list the jms event we listen to. The graphManagerEventListener extension point allows to contribute a class that would be called when such event are received.

The default implementation is DocumentVersionnedGraphManager. It listens to the documentCheckedIn event. This event is sent when a new UUID is created for a document. The manager add, to each annotation that annotates the old UID, a new annotates statement with the new UID. The result is that when you version a document, the old version will have the same annotation than the new version, but annotation on the new version will not be seen in the old version.

51.4. NXAS Web

51.4.1. Overview

The NXAS HTML Client is the a web interface that can be used by the end user. It is written in GWT. It can be configured in two ways. The "normal" Nuxeo way of extension point, or adding attribute or object to the html page to be read by the gwt module.

Implementation Note

51.4.2. Extension Point

The WebAnnotationConfigurationService allows to contribute:

- the different types of annotations you want to manage in your application. The default contribution provides the 6 annotea protocols types: SeeAlso, Question, Explanation, Example, Comment, Change and Advice.
- the filters you may want to use in your application. They are displayed as button on the left side. The default contribution shows no filter.
- the displayed fields: fields you want to display in the annotations list in addition to the icon and the date.
- a class implementing the `UserInfoMapper` interface to add to the `WebConfiguration` some informations about the current user.
- a class implementing the `WebPermission` interface to let know to the client if annotations are allowed or not on the current document.

For more informations and examples, see the files `web-configuration-service-contrib.xml` and `web-configuration-service-framework.xml` located in the `nuxeo-platform-annotations-web` project.

51.4.3. Configuration

51.4.3.1. Panel Element

The client uses a div with id "display" to add the annotation application to the current page.

51.4.3.2. Preview setting

To allow integration with the preview module, additional configuration is possible via a div with an id of "preview-setting". Each configuration is an attribute of the div.

- imageOnly value should be true or false. If only the image should be annotated in the frame.
- multipleImageAnnotation should be true or false. If more than one image should be annotated.
- The annotation module uses 2 functions to transform the xpointer gotten from clicking on an image, and to transform a xpointer to a xpointer that make sens in the previewed image. It get the function from top[xPointerFilterPath] and top [pointerAdapter]. xPointerFilterPath and pointerAdapter are the value of the same named attribute.

51.5. Annotation Service Facade

The facade provides the integration of the service into JEE infrastructure. EJB3 Stateless Session Bean facade on top of the runtime service. This makes the Annotation Service remotable and accessible via the standard Nuxeo Lookup facility :

```
Framework.getService(AnnotationService.class)
```

51.6. Annotation Service HTTP Gateway

51.6.1. Overview

The Annotation Service HTTP gateway provides the http API on top of the java APi of the Annotation Service It is a servlet that implements the annotea protocol.

51.7. References

- W3C Annotea Protocol
- *Nuxeo preview addon*
- *XPointer*

Chapter 52. Virtual Navigation

52.1. Virtual Navigation

This section is speaking about a Nuxeo addon that provides a new way to navigate in documents.

52.1.1. Virtual Navigation presentation

Virtual navigation is opposed to the physical navigation. The physical navigation is the intuitive way to browse your documents as you created them, folders contain documents. Virtual Navigation is based on meta-data linked with every documents. In Nuxeo 5 each document has a set of meta-data that make him more precise and rich. The Virtual Navigation can provide a navigation tree built on meta-data and more precisely on every vocabulary based data.

Now the basic Virtual Navigation configuration offers a navigation through coverage and through subjects. For each document it is possible to determine a country that is relative to the document and one or many subjects that corresponds well to it.

To activate the Virtual Navigation you have to add the Virtual Navigation addon in the plugin extension folder of your JBoss server. After restarting the server you can see a new widget in the left hand corner of Nuxeo 5. You can switch between navigation style easily.

The standard physical navigation is still here and the two Virtual Navigation are selectable on the right. As you can see the Virtual Navigation keeps a folder based navigation but this tree browsing is built from vocabularies defined in your Nuxeo 5 instance. Coverage and Subjects are 2 vocabularies that exist in the basic Nuxeo 5 configuration. You may know that tree based vocabularies are not limited in depth and you can have 2 or more level in your vocabulary, for example Coverage allows to select a **continent** and a **country**.

To have good results with Virtual Navigation you have to fill meta-data linked to documents. The Meta-Data view allows the user to provide his own meta-data. Select a country and/or subjects to make your Virtual Navigation efficient. See below to view the Meta-Data configuration screen.

When selecting a node in a Virtual Navigation tree it will execute a request to find every documents that contain the wanted meta-data. See the screen below for an example: if you select the Art/Archiecture couple in the Subjects navigation tree, every documents that contains Art/Architecture in their Subjects meta-data will be displayed on screen.

52.1.2. Virtual Navigation configuration

The basic Virtual Navigation configuration, as seen above, offers 2 way to navigate through meta-data in addition to the classic physical navigation, by coverage and by subjects. Here is a complete walkthrough to add a new virtual navigation to a customized project.

52.1.2.1. Required precondition

You have to create a vocabulary that will be used for the new Virtual Navigation, keep in mind that you can set up a multi-level vocabulary, so you can imagine a vocabulary with parent and children. There is no limitation on this side.

You have to select the element in a schema that will store your vocabulary data. In the basic Nuxeo 5 configuration it is dublincore:coverage and dublincore:subjects that store vocabularies data for Virtual Navigation. If fact, when you select a country (coverage) in a document's meta-data, the country is stored in the dublincore:coverage element. It's not mandatory to select a dublincore element, you can select an element coming from your own schemas.

52.1.2.2. Set up your vocabulary

If you have a "only-one-level" vocabulary you can create a file to store it like this :

Example 52.1. Example to create your own vocabulary

```
id, label, obsolete
"art","label.directories.topic.art","0"
"human sciences","label.directories.topic.humanscience","0"
"society","label.directories.topic.society","0"
"daily life","label.directories.topic.dailylife","0"
"technology","label.directories.topic.technology","0"
```

If you have a "multi-level" vocabulary you have to split him in sub vocabulary, and keep in mind that you have to indicate which vocabulary element is the child of which parent vocabulary element, here is an example, the full example can be found in Nuxeo 5 sources, check topic.csv and subtopic.csv:

Example 52.2. Example to create your own multi-level vocabulary

```
CONTENT OF THE FIRST FILE (parents)

id, label, obsolete
"art","label.directories.topic.art","0"
"human sciences","label.directories.topic.humanscience","0"
"society","label.directories.topic.society","0"
"daily life","label.directories.topic.dailylife","0"
"technology","label.directories.topic.technology","0"


CONTENT OF THE SECOND FILE (with a "parent" parameter added)

id, label, parent, obsolete
"architecture","label.directories.subtopic.architecture","art","0"
"comics","label.directories.subtopic.comics","art","0"
"rights","label.directories.subtopic.rights","human sciences","0"
"economy","label.directories.subtopic.economy","human sciences","0"
...
```

When you have created your different files for vocabularies you have to register them in an extension point as new vocabularies. In an xml file that you won't forget to place in a OSGI-INF directory and won't forget to register in the MANIFEST.MF file of the package, you have to contribute the following extension point.

For a full example you can check the **nxdirectories-contrib.xml** file of the **webapp-core** package. Don't forget to use the **xvocabulary** schema if your vocabulary is a child of another, if the vocabulary is the first parent or is alone just use the **vocabulary** schema. Don't forget to indicate the parent if your vocabulary has more than one level with the following tag **<parentDirectory></parentDirectory>.**

Example 52.3. How to register new vocabularies

```
<extension target="org.nuxeo.ecm.directory.sql.SQLDirectoryFactory"
  point="directories">

//here your new vocabularies contribution

</extension>
```

52.1.2.3. Set up a new Document Type for search purpose

The query based search service of Nuxeo 5 requires that you create a document type that will be a base for a

document model to register data handled by the query. To understand it more here is an example:

You are browsing your documents by coverage. You are selecting the path Europe/France in the tree. The data Europe/France is the base of the query and need to be registered in a document model created from a document type. In this case the document type can be very simple, cause the query must register only one data at a time. Creating a document type with only one schema that contains one field will be enough.

This document type will be referenced as **query document type** in this walkthrough. Create it and register it as a normal document type. See Nuxeo Book part 6 to get more informations about document type creation.

52.1.2.4. Set up new navigation tree

Now you must contribute to another extension point to create a new navigation tree based on vocabularies you contributed just before. See below for an example of the file you have to create. For a full example you can see the **directorytreemanager-contrib.xml** file in the **virtualNavigation** package.

The new tree contribution needs many informations, **a queryModel** to indicate the query you will use to get your documents, a **schema** and a **field** coming from the query document type you just set up in part 3, an **outcome** that indicates the page where documents will be displayed after the request and a list of vocabularies you can indicate with the tag **<directory></directory>.**

Example 52.4. How to register new navigation trees

```
<require>
  org.nuxeo.ecm.webapp.directory.DirectoryTreeService
</require>

<extension
  target="org.nuxeo.ecm.webapp.directory.DirectoryTreeService"
  point="trees">

//here your new navigation tree contribution

</extension>
```

52.1.2.5. Set up the queryModel and the Result Provider

Now you need to contribute to 2 extension points that will set up the query and the results provider. Here are the 2 extension point you will need to contribute. For a full example see the file **querymodel-contrib.xml** and **resultsprovider-contrib.xml** in the **virtualNavigation** package.

The query model contribution needs many information. In the **docType** parameter you have to put the name of the documentType you created in part 3. In the <predicate></predicate> tag you must set up the name of the element (field) that is the field storing the data used for virtual tree construction.

EXAMPLE :

You want to browse documents by coverage, each of your document have a coverage registered in dc:coverage field. You have to use the **dc:coverage** parameter.

If there is no prefix set the name like this **schema:field**, if there is a prefix set it up like this **prefix:field**. In the operator param put the value STARTSWITH. In the <field/> tag put the schema and the name of the field coming from the **query document type** you set up in part 3.

Example 52.5. How to register query model and results provider

```
<extension
   target="org.nuxeo.ecm.core.search.api.client.querymodel.QueryModelService"
   point="model">

//Here your new query model
```

```
</extension>


<extension
    target="org.nuxeo.ecm.webapp.pagination.ResultsProviderService"
    point="model">

//Here your new result provider

</extension>
```

52.1.2.6. Set up the search service

Now you need to register a new indexing configuration in the search service. You will index the field that register the data you want to apply the search on, it is the same data as saw at the end of part 5 . Here is the extension point you have to contribute. For a full example see the **nxsearch-contrib.xml** file in the **search-core** package.

You will have to create a new <resource></resource> tag if it doesn't exist for your schema. For the **name** parameter you have to put the name of your schema, **type** param must be "schema" and **indexAllFields** param must be "true". In a <field/> tag you have to indicate the indexing strategy, the **type** param must be "**Path**", if the field that register your vocabulary is a complexType (list of String for example) you can add the parameter **multiple** and set it at "true".

Example 52.6. How to register the search configuration

```
<extension
    target="org.nuxeo.ecm.core.search.service.SearchServiceImpl"
    point="resource">

//Here your new search configuration

</extension>
```

52.1.2.7. Other details

You must not forget to add some navigation cases in a deployment-frgament.xml file. Those navigation cases must correspond at each **outcome** you set up in part 3. Each page named in navigation cases shall display search result so you have to customize those pages with the good queryModel call. For a full example see **coverage_virtual_navigation.xml** .

All theme contributions and tree navigation widget already exist in **webapp-core** and **virtualNavigation** package, you can take a look for information purpose.

Indeed you have to add vitualNavigation to Nuxeo 5 plugin if you want a fully fonctional Virtual Navigation immediatly.

Chapter 53. Metadata Extraction Service

53.1. Introduction

There are cases when some informations could be retrieved from attached files and keep this info as regular Document properties. This info is refered to as metadata as it is a descriptive info refering the document (file) from which it is extracted. For example given a MS Word document (file) we could retrieve info (metadata) like *author, title, creation date* etc, simply by reading document headers with an appropriate library (that knows how to parse and keep a MS Word document internal structure).

53.2. Metadata extraction module

The metadata extraction feature is composed of several projects:

- `nuxeo-platform-metadataext-api`

- `nuxeo-platform-metadataext-core`

- `nuxeo-platform-metadataext-facade`

- `nuxeo-platform-metadataext-plugins`

The core part defines the metadata extraction component with a `MetaDataExtractionManager` implementation: `org.nuxeo.ecm.platform.metadataext.services.MetaDataExtractionService`.

The service is normally invoked by a dedicated `CoreListener` that passes a `DocumentModel` for which an extraction could be defined. Also the metadata could be extracted by direct invocation from a client code through existing EJB3 facade.

The plugin part provides a plugin (more to come) which is specific to the MS-Word document type. The plugin is a Transformation Plugin (as defined by the Transformation service) and gets invoked as part of a defined transformation chain. The apropriate transformation will be called by the metadata extraction service if there is a contribution for the given Document type, etc.

53.2.1. Defining a contribution for metadata extraction

A contribution can be defined by adding an XML file with the following structure:

```
<extension
    target="org.nuxeo.ecm.platform.metadataext.services.MetaDataExtraction"
    point="extractions">

  <meta-data-extraction inputField="file:content"
      transformationName="MSWordMDExt">

    <outputParams>
      <param propertyName="dc:title">title</param>
      <param propertyName="dc:description">comments</param>
      <param propertyName="dc:created">creationDate</param>
      <param propertyName="dc:modified">creationDate</param>
      <param propertyName="dc:contributors">authors</param>
    </outputParams>

    <coreEvent>documentCreated</coreEvent>
    <coreEvent>documentModified</coreEvent>

    <docType>File</docType>

  </meta-data-extraction>

</extension>
```

The important aspects for MetaDataExtraction service are:

- specification of a source (blob) from where metadata will be extracted. This is defined by the value *`inputField`* which should be a blob document property.

- the mapping of output parameters. This defines a corespondence between the map entries returned as the result of transformation (extraction) and the Document properties names to which the results will be written back.

- the list of core events for which the extraction will be performed.

- the list of document types for which the extraction can be applied.

53.2.2. Specifying input parameters

If the extraction is requiring additional information aside from the provided blob, we can add input parameters to an extraction definition like this:

```
<inputParams>
  <param propertyName="dc:title">title</param>
</inputParams>
```

A map with input parameters will be passed to the extraction plugin having the key the value of tag `param` (in this case 'title').

53.2.3. Chaining extractions

In extreme cases the extractions of metadata could be performed in several steps. That is using the extracted parameters in phase as input parameters for the next phase. This could be achieved by specifying the input params whose values could have been set by a previous extraction (transformation).

53.2.4. Creating a plugin for metadata extraction

We first define a class implementing `org.nuxeo.ecm.platform.transform.interfaces.Plugin` interface like:

```
public class MSWordMDExtractorPlugin extends AbstractPlugin {

    @Override
    public List<TransformDocument> transform(Map<String, Serializable> options,
            TransformDocument... sources) throws Exception {
      ...
    }
...

}
```

Overriding the `transform` method is a must and in this case we will have only one `TransformDocument` as source.

The blob from which metadata could be extracted can be retrieved from `TransformDocument` like: `Blob srcBlob = sources[0].getBlob()`.

After the useful information is extracted the properties should be set to a `TransformDocument` that will be returned:

```
TransformDocumentImpl res = new TransformDocumentImpl();
res.setPropertyValue("title", extractedTitle);
```

53.2.5. Using a metadata extraction plugin

Having a plugin class defined that knows how to extract information from a specific type of file (MS Word document, MS Excel, PDF, OO doc etc) we can add the contributions necessary for the metadata extraction to

take place.

We need to define first the transformation plugin and the transformation contributions. After that the defined transformation can be refered in the metadata extraction specific contribution.

Step 1: define the transformation plugin:

```
<extension
    target="org.nuxeo.ecm.platform.transform.service.TransformService"
    point="plugins">

  <documentation>
    Set of default transformation plugins for metadata extraction.
  </documentation>

  <plugin name="MSWordMDExtPlugin"
      class="org.nuxeo.ecm.platform.metadataext.plugins.MSWordMDExtractorPlugin"
      destinationMimeType="application/msword">
    <sourceMimeType>application/msword</sourceMimeType>
  </plugin>
</extension>
```

Step 2: define the transformation:

```
<extension
    target="org.nuxeo.ecm.platform.transform.service.TransformService"
    point="transformers">

  <documentation>
    Set of default transformation chains for metadata extraction.
  </documentation>

  <transformer name="MSWordMDExt"
      class="org.nuxeo.ecm.platform.transform.transformer.TransformerImpl">
    <plugins>
      <plugin name="MSWordMDExtPlugin" />
    </plugins>
  </transformer>
</extension>
```

Step 3: define the metadata extraction specific contribution

```
<extension
    target="org.nuxeo.ecm.platform.metadataext.services.MetaDataExtraction"
    point="extractions">

  <meta-data-extraction inputField="file:content"
      transformationName="MSWordMDExt">

...
```

This last one is refering to the above defined transformation.

Chapter 54. Unicity Service

The Unicity service sends a new Message on the JMS bus every time you upload a file already on the server.

54.1. How does it works ?

When creating or updating a file, a digest is computed using the whole file and an encoding algorithm. Next step is an nxql query wich gets every Document with the same Digest. If such documents exist, there references will be forwarded on JMS bus.

54.2. What you need:

The message's eventID is *duplicatedFile*. You need it to catch the event. The documentLocation list is available in the info map of the CoreEvent using *duplicatedDocLocation* key.

54.3. Configuration

How to configure unicity extension point. TODO: needs to be made clearer.

```
<enabled>
Default value is false. Use true if you want to enable this service.
</enabled>
<algo>
Default encoding algorithm is sha-256. You can choose every algorithm supported by java.security.MessageDigest.
</algo>
<field>
A field is an xpath expression giving a particular field of your schema.
It's the reference of a file on the server.
The type's field must be nxs:content.
You can use as many field as you want.
</field>
```

54.4. Snippets

Some code to get you started.

Example 54.1. Sample unicity extension point contribution to the FileManager service.

```
<?xml version="1.0"?>
<component name="org.nuxeo.ecm.platform.filemanager.service.FileManagerService.Plugins">
  <extension
      target="org.nuxeo.ecm.platform.filemanager.service.FileManagerService"
      point="unicity">
    <unicitySettings>
      <enabled>true</enabled>
      <algo>sha-256</algo>
      <field>content</field>
    </unicitySettings>
  </extension>
</component>
```

Example 54.2. Sample code to retrieve DocumentLocations.

```
public void onMessage(Message message) {
    try {
        Serializable obj = ((ObjectMessage) message).getObject();
        if (!(obj instanceof DocumentMessage)) {
```

```
            return;
        }
        DocumentMessage doc = (DocumentMessage) obj;

        String eventId = doc.getEventId();

        if ("duplicatedFile".equals(eventId)){
            Object[] documentLocations = (Object[]) doc.getEventInfo().get("duplicatedDocLocation");
            for (Object documentLocation: documentLocations) {
                log.debug(((DocumentLocation)documentLocation).getDocRef());
            }
        }
    }
}
```

Chapter 55. Nuxeo Mail Service

55.1. Presentation

Nuxeo mail service enables users to send and receive mail. It also makes it possible to apply a serie of actions on messages located in folders.

The service is a wrapper around JavaMail . Knowledge of this API is necessary for advance use of the service.

55.2. Basic Use

55.2.1. Configuration

Use of the mail service is done via session. You need to configure a session factory that will be able to find the correct server parameter to connect to. Then, the session can use the mail server. The simplest scenario is to have mails sent and received using the credential of one user (see property fetcher for more complex use).

To pass the server parameter in your extension point, use a simple properties fetcher. For a server using TLS connection with SMTP and IMAP, a configuration could look like:

```
<extension target="org.nuxeo.ecm.platform.MailService"
   point="sessionFactory">
   <sessionFactory name="mysession" fetcherName="simple">
     <properties>
       <property name="mail.store.protocol">imap</property>
       <property name="mail.transport.protocol">smtp</property>
       <property name="mail.smtp.port">587</property>
       <property name="mail.host">mail.mycompany.com</property>
       <property name="mail.smtp.host">mail.mycompany.com</property>
       <property name="mail.imap.starttls.enable">true</property>
       <property name="mail.imap.ssl.protocols">TLS</property>
       <property name="mail.smtp.ssl.protocols">TLS</property>
       <property name="mail.user">mailuser@mycompany.com</property>
       <property name="mail.from">mailuserd@mycompany.com</property>
       <property name="mail.imap.port">143</property>
       <property name="password">myuserpassword</property>
       <property name="user">maiuser@mycompany.com</property>
     </properties>
   </sessionFactory>
 </extension>
```

With such a configuration you can send mails using the mail service, passing the text of the mail as a String.

```
MailService mailService = Framework.getService(MailService.class);
mailService.sendMail("My interesting mail.", "The subject of the mail",
        "mysession", new Address {internetAddress});
```

To process mails in mail folders, you need to create a pipe and add actions to it. Each action in the pipe will be applied to each mail in the folder. A typical pipe will include action to set the delete flag so the mail is deleted when the pipe exits, transforms the mail, stores it, and maybe answers it. For all those actions to communicate with each other, you can pass objects into the context. Context is passed from action to action.

The package `org.nuxeo.ecm.platform.mail.action` contains a set of Action classes you can use or extend. Note that some actions expect objects to be in the map.

A MessageAction is an action done on a message:

```
public interface MessageAction {
   boolean execute(ExecutionContext context) throws Exception;
   void reset(ExecutionContext context) throws Exception;
}
```

```
Message message = context.getMessage();
String infoFromPreviousAction = context.get("info.from.previous.action");
context.put("info.for.next.action", "don't save this message");
```

If the returned value of an action is false, then the processing of this message is stopped and the next message will be processed.

The MailBoxAction enables using those service in a simple way, a typical use would be:

```
MailBoxActions mba = mailService.getMailBoxActions("mySimpleFactory", "INBOX");
mba.addAction(new ConvertToDocumentModelAction());
mba.addAction(new StoreDocumentModelAction());
mba.addAction(new SendMailReplyAction());
mba.execute();
```

55.3. Advanced Use

The Mail Service provides access to JavaMail object for advanced use:

```
Store getConnectedStore(String name) throws MessagingException;
Store getConnectedStore(String name, Map<String, Object> context) throws MessagingException;
Transport getConnectedTransport(String name) throws MessagingException;
Transport getConnectedTransport(String name, Map<String, Object> context) throws MessagingException;
Session getSession(String name);
Session getSession(String name, Map<String, Object> context);
```

The mail service comes with a simple fetcher. The simple fetcher enables passing the configuration for the mail server using key/value pair. If you need more complex configuration, such as having specific properties for each user fetched from a directory, you can contribute a new fetcher using the propertiesFetcher extension point.

```
<extension target="org.nuxeo.ecm.platform.MailService"
    point="propertiesFetcher">
  <propertiesFetcher name="simple"
      class="org.nuxeo.ecm.platform.mail.fetcher.SimplePropertiesFetcher"/>
</extension>
```

To test your mail service, jes mail server is a simple, light weight mail server you can use locally. You can either start it manually or use the `nuxeo-platform-mail-test` bundle to do it.

Chapter 56. Imaging

56.1. Introduction

Nuxeo Platform Imaging provides picture management to Nuxeo Web Platform and RCP. It offers minor picture transformation like rotation or resizing.

There are several Imaging addons for the Nuxeo platforms. They are listed in the following sections.

56.2. Imaging API

This addon provides a pluggable adapter for picture manipulation. It makes it possible to manipulate images resources that are not documents (in the document type sense) themselves. Concretely, it makes it possible to still use the Seam component that deals with images for some processings (for example rotation and resizing) but adapted to a specific schema.

56.3. Imaging transform

This addon is where the available transformations are stored. Available transformations are rotation, resizing and cropping for the following image types: JPEG, GIF, BMP, TIFF. PNG is sadly not supported yet.

56.4. Imaging web

This addon provides Picture and PictureBook document types as well as drag'n drop plugin, slideshows or thumbnails.

56.5. Imaging core

This core addon provides services for common picture processing operation, metadata extraction and mime-type detection. With this addon, one can choose the picture processing library to use through the library selector service. Right now the available libraries are Mistral and ImageJ.

Chapter 57. Nuxeo Preview Addon

The `nuxeo-platform-preview` addon provides services, transformers and adapters to generate HTML previews from of a Nuxeo DocumentModel.

57.1. Overview

The preview addon is composed of several layers:

- UI part

  Adds a Preview tab that displays the preview inside an IFRAME.

- Transformer part

  The Preview addon contributes transformers that are dedicated to generating HTML out of most file formats.

- Service and Adapters

  The Preview addon includes services that allow you to define `PreviewAdapters` for each type of Document. These adapters are responsible for defining how the preview will be generated:

  - define the field(s) used for preview

    Because a Nuxeo document contains a lot of differents fields, the adapter must determine what are the default fields used to generate the preview.

  - define how preview is generated

    Basically preview can be generated using some transformers or using some fields that already contains pre-generated HTML preview

  The services provided by the Preview addon let you configure the `PreviewAdapter` depending on your document types and manage some cache for transformer based previews.

57.2. Installing the Preview Addon

Installing the Preview addon is fairly simple.

57.2.1. Deploy the Addon

Just drop the `nuxeo-platform-preview-5.X.jar` in the `nuxeo.ear/plugins/` directory and restart your server.

57.2.2. Configure Transformers

The preview service relies on HTML transformers that uses external tools to achieve the transformation. You need to have them installed and configured before the preview addon will work.

The preview service comes with two HTML transformers:

- `PDF2Html`

  This transformer generates HTML from PDF files.

  This transformer relies on the command line tool `pdftohtml` This tool can be found as package for most

linux distrib (for instance, included in `poppler-utils` in Ubuntu). This tool can also be installed on MS Windows platform, see http://sourceforge.net/projects/pdftohtml/

Once you have installed `pdftohtml` on your server, you must configure Nuxeo to let him know where is the command line tool. For this, you can use a extension point to define what is the temporary directory to use and where is `pdftohtml` command.

Default settings should be ok for most Linuxes, but you will have to do the configuration if you use a MS Windows box.

To do this configuration, create a file named `pdftohtml-config.xml` in the `nuxeo.ear/config` directory

```
           <?xml version="1.0"?>
<component name="myproject.pdftohtml.config.contrib">
  <extension target="org.nuxeo.ecm.platform.preview.transformers.CLTransformerPluginParameterManagerComponent"
    point="cltpParameters">

  <CLTranformerPluginParameters name="Pdf2HtmlParams" targetTransformerPlugin="Pdf2Html">
    <parameters>
      <!-- default Linux parameters
      <parameter name="commandString">/usr/bin/pdftohtml</parameter>
      <parameter name="tmpDir">/tmp/</parameter> -->
      <!-- MS Windows parameters -->
      <parameter name="commandString">C:\\Program files\\pdftohtml\\pdftohtml.exe</parameter>
      <parameter name="tmpDir">C:\\Temp\\</parameter>
    </parameters>
  </CLTranformerPluginParameters>
  </extension>
</component>
```

- `Any2Html`

  This Transformer is in fact a chain using `Any2pdf` and `PDF2html`. This means you need to have `PDF2html` working but also OpenOffice.org in listen mode (this is documented earlier in this book).

57.3. Extensions and Pluggability

The Preview addon provides several extension solutions.

57.3.1. Pluggable Adapters

Preview AddOn uses the standard `DocumentModelAdpater` system. For that it defines a `HtmlPreviewAdapter` interface that must be implemented by each adapters. In order to let you choose the Adapter implementation according to the Document type, the `DocumentAdapterFactory` used for the `HtmlPreviewAdapter` is pluggable. This factory uses the `PreviewAdapterManager` service to determine what implementation of the `HtmlPreviewAdapter` should be returned when you use

```
doc.getAdapter(HtmlPreviewAdapter.class)
```

.

Defining custom preview adapter is done via a dedicated extension point: you register a Factory for a given DocumentType. This factory will be used to create the implementation of the `HtmlPreviewAdapter` from the DocumentModel. Here is example of such a contribution:

```
         <?xml version="1.0"?>
<component name="org.nuxeo.ecm.platform.preview.adapter.contrib">
  <extension target="org.nuxeo.ecm.platform.preview.adapter.PreviewAdapterManagerComponent"
    point="AdapterFactory">

    <previewAdapter name="notePreviewAdapter" enabled="true">
      <typeName>Note</typeName>
      <class>org.nuxeo.ecm.platform.preview.adapter.factories.NotePreviewAdapter</class>
    </previewAdapter>
```

```
  </extension>
</component>
```

This contribution defines the Adapter factory for the Note Document type. Contributed factories must implement the `PreviewAdapterFactory` interface (that contains only one method!).

The implementation returned by the factory can be based on one of the 2 base classes provided inside the addon.

- `TransformerBasedHtmlPreviewAdapter`

  This base class provides all the build in features to create a preview adapter that uses the Transformation Services to generate the html preview.

  Default usage in Nuxeo is the generic adapter for the all the documents that contains the `file` schema.

- `PreprocessedHtmlPreviewAdapter`

  This base class provides all the build in features to create a preview adapter that uses pre-processed HTML preview that is stored inside the document.

  Default usage in Nuxeo is the adapter for the Note Document type.

  For most usages, you just be able to contribute a factory that create a new `PreprocessedHtmlPreviewAdapter` with the good arguments.

57.3.2. Pluggable HTML Transformers

The Preview addon contributes some command line based transformers.

It also provides you with a base class `AbstractCommandLineBasedTransformer` that can be used to easily make another CommandLine based Transformer.

We also provide a `CLTransformerPluginParameterManager` service that is used by the base class to extact the configurable parameters.

57.4. Previews and URLs

The preview addon includes a restlet.

This restlet provides a RESTful access to the preview service. This means that you can retrieve the preview of a document by just using a simple GET URL.

Preview URLs have this form :

```
http://{server}:{port}/nuxeo/restAPI/preview/{document_uuid}/{previewfield}/
```

where:

- `document_uuid`

  is the the document uuid that is already present in all Nuxeo urls

- `previewfield`

  is the xpath of the field that should be used as main file for the preview

  Depending on the underlying adapter implementation, this field may or may not be relevant. Use 'default' to let the adapter implementation choose the right field for you.

Please note that the last / is important. As HTML preview can contain nested images, the base URL must end with a /. The Restlet system handles a cache to avoid refetching the preview data at each call.

Chapter 58. Nuxeo Tiling service Addon

The `nuxeo-platform-imaging-tiling` addon provides a service and adapters to generate picture tiles from a DocumentModel containing a picture or from a blob.

Tiling picture is usefull when you want to display in a web browser a very big image : accoring to the display settings and to the zoom level, the picture tiling service will generate you the needed picture tiles.

58.1. Overview

The Tiling Service addOn is composed of several layers :

- REST API part

  Provides a rest API to get the picture tiles from the server.

  Depending of the URL, this API could be used to get the actual tiles (ie: picture), or to get information about the picture and the tile grid (XML or JSON).

- Adapter part

  An DocumentModelAdapter is provided to adapt the DocumentModel interface to an interface dedicated to picture tiling.

  Default adapter simply extract the blob from the DocumentModel using either the default file schema, or using a provided xPath.

  Other adapters can be provided to implement specific logic for source picture retrieveal (like multi-view-pictures) or even extract pre-computed tiles from the DB.

- Tiling Service

  This service is responsible for :

  - managing tiles cache

    Because picture tiling can be long, this service maintains a FileSystem based cache of the input pictures, the generated tiles, and also some temporary resources used by the tiling process.

    The service also manage a configurable garbage collector to clean up this cache.

  - manage PictureTilers

    A PictureTiler is a Java Class that can extract a tile from a picture.

    Current implementation provides ImageMagick based and Gimp based PictureTiler.

  Default configuration requires ImageMagick :

  - because it is faster
  - because some treatments can not be done via gimp
  - because setup is easier

58.2. How tiles are defined and computed

In order to define a tiling you will need to define 4 input parameters:

- The input picture.

  This input is represented by the `ImageResource` className in the API.

  This `ImageResource` is usually provided via a DocumentModel adapter that will encapsulate the logic to find the input picture and will create a `DocumentImageResource`.

  You may also directly create a `BlobResource` from a simple Blob.

- The tiles width in pixel
- The tiles height in pixel
- The maximum number of tiles X/Y

  This parameter represents the maximum number of tiles on the X or Y axis, that are needed to display the complete picture.

  Basically it represent the grid size you will need to display the complete image using this tiling.

As you may have noticed the zoom factor is not an input parameter. In fact it is an output parameter : the service will compute the zoom factor that correspond to the tiling definition.

This choice may seem strange, but in most of the case, this is easier to use : you don't always know the size of the underlying image, but you know the space you have to display the tiled image (total width/height and number of tiles) and you know that you want to start with an image that is fully displayed.

Here is a quick explanation on how the tiling service work.

- compute needed picture size.

  From the tiling grid definition, the service will compute the needed input picture size.

- find a scaled down picture if available

  When the needed picture size is really smaller that the real input image, it may be interesting to generate a scalled down image.

  This will slow down the process for the computation of the first tile, but will speed up the computation of all other tiles.

  This is particularly true for format like JPEG that support to easily extract from them a smaller image.

  This speed up the reading part (less bytes to read from DB/Hard Drive) and speed up the final scalling.

  When scaling down the input picture is possible, the resulting image will be stored into the cache.

- select the tiler
- let the tiler do it's job

  Depending on the tiler implementation, it may generate one tile at a tile (default for ImageMagick based tiler) or 9 tiles at a time (default for gimp tiler) or all tiles (possible for gimp tiler).

  This behavior is transparent for the caller since, all generated tiles are cached.

58.3. Installing the picture tiling Addon

58.3.1. Requirements

ImageMagick 6.3.7 or later is needed : the Tiling Service uses the ImageMagick stream command that is not available in 6.2 (associated packages are libmagick and imagemagick).

Current implementation has only be tested under Linux.

If you want to use the Gimp tiler, you will need Gimp 2.x with python extensions.

58.3.2. Deploy the addon

Just drop the `nuxeo-platform-imaging-tiling-XXX.jar` in the `nuxeo.ear/plugins/` directory and restart your server.

58.3.3. Configuration

Configuration can be done using a extension point. Just create a file called pictures-tiles-config.xml in nuxeo.ear/config.

Use this extension point contrib to :

- define the imagemagick command path (if default is not ok)
- define the directory that will be used for cache
- define the cachesize
- define the GC parameters

NB : default config should be ok for any linux based system where imagemagick is setup via the package manager

Here is an example of a configuration file :

```
<?xml version="1.0"?>
<component name="my.projects.tiles.config">
  <require>org.nuxeo.ecm.platform.pictures.tiles.default.config</require>
  <extension target="org.nuxeo.ecm.platform.pictures.tiles.service.PictureTilingComponent"
    point="environment">
  <environment>
    <parameters>
      <!-- Gimp path variables -->
      <parameter name="GimpExecutable">gimp</parameter>
      <!-- ImageMagick path variables -->
      <parameter name="IMConvert">convert</parameter>
      <parameter name="IMIdentify">identify</parameter>
      <parameter name="IMStream">stream</parameter>
      <!-- global env variables -->
      <parameter name="WorkingDirPath">/tmp/</parameter>
      <!-- Max Disk cache usage in KB -->
      <parameter name="MaxDiskSpaceUsageForCache">50000</parameter>
      <!-- GC Interval in Minutes -->
      <parameter name="GCInterval">10</parameter>
    </parameters>
  </environment>
  </extension>
</component>
```

58.4. Testing the tiling service

The tiling service can be very easily tested.

58.4.1. URL and restAPI

The Rest API for the tiling service use the felowing url pattern :

```
http://{server}/nuxeo/restAPI/getTiles/{repoId}/{docUUID}/{tWidth}/{tHeight}/{maxTiles}?format={format}&x={x}&y=
```

where

- {server}

  is the server DNS name of IP

- {repoId}

  is the id of the core repository that is used (use default if you don't know what it is).

- {docUUI}

  The uuid of the document that contains the image. If you use the default adapter, this document must use the `file`schema and contain an image.

- {tWidth}

  the width in pixel of the tiles.

- {tHeight}

  the height in pixel of the tiles.

- {maxTiles}

  maximum number of tiles X/Y for the full image

- {format}

  the format of the tile information (JSON/XML)

  this is used only if x and y are not supplied

- {x}

  x of the tile to be generated (starting from 0)

- {y}

  y of the tile to be generated (starting from 0)

Here are some examples :

```
http://server/nuxeo/restAPI/getTiles/default/950b0d27-2ca4-43e4-bb12-598ad6d64e86/200/150/2
```

Will send you the tiling information in XML for the picture contained in the doc 950b0d27-2ca4-43e4-bb12-598ad6d64e86 with a tile size of 200x150x2.

```
http://server/nuxeo/restAPI/getTiles/default/950b0d27-2ca4-43e4-bb12-598ad6d64e86/200/150/2?format=JSON
```

Will send you the tiling information in JSON for the picture contained in the doc 950b0d27-2ca4-43e4-bb12-598ad6d64e86 with a tile size of 200x150x2.

```
http://server/nuxeo/restAPI/getTiles/default/950b0d27-2ca4-43e4-bb12-598ad6d64e86/200/150/2?x=1&y=1
```

Will send you the tile (1,1) for the picture contained in the doc 950b0d27-2ca4-43e4-bb12-598ad6d64e86 with a tile size of 200x150x2.

58.4.2. simple test client

The RestAPI provide a very basic test client. This test client was used during the development : this client is not pretty, not IE compliant, not dummy user proof.

You can access this test client via the rest API, just add the test=true to the URL.

```
http://server/nuxeo/restAPI/getTiles/default/950b0d27-2ca4-43e4-bb12-598ad6d64e86/200/150/2?test=true
```

Chapter 59. Nuxeo DocumentLink Addon

The `nuxeo-platform-documentLink` addon provides services, and adapters to handle "application level" proxies pointing to DocumentModels.

59.1. Overview

There are several use cases that require having the same document available from several places. Typical use cases include publication (same document visible in multiple sections), personnal workspaces (user want to have some document into his workspace without copying them) ... Nuxeo Core provides a build-in feature called Proxies that can be used in most cases. The current limitations of proxy system in Nuxeo Core 1.4.x include :

- A proxy can only point to a checked-in DocumentModel (ie : a version)
- A proxy is always totally equivalent to the target `DocumentModel` (ie : same schemas and field values)

DocumentLink provides a proxy system implementation on top of the Core using the `DocumentModel` adapter system. DocumentLink extends the Core proxy system to provide some additionnal features :

- A DocumentLink can point to a checked-in or checked-out DocumentModel
- A DocumentLink can have it's own schemas and fields.
- A DocumentLink can mask some of the schemas/fields of the target DocumentModel

As an example, if you have a DocumentModel DocA with title "Document A" and description "description A", you can create a DocumentLink DLA that will point to DocA and have title "DocLink A" but will always return the description contained in DocA.

59.2. How does it work

The DocumentLink system uses DocumentModelAdapter to adapt the default `DocumentModel` implementation to a specific implementation that handles the logic for dispatching attributes access across the DocumentLink and the target Document.

The DocumentLink adapter (`DocumentLinkAdapter`)implements the `DocumentModel` interface and will by default have the felowing behavior :

- get internal `DocumentModel` value if the target schema/field is available.
- return the value stored inside the target `DocumentModel` otherwise.

The reference of the target `DocumentModel` is stored in a dedicated schemas (named documentLink).

This specific schema contains :

- target DocumentRef
- target DocumentRepository
- list of schema that will never be masked by the DocumentLink

  This means that for this schema the DocumentLink will always return the values stored in the target `DocumentModel` even if the DocumentLink it self has these schemas.
- list of fields xPaths that will never be masked by the DocumentLink

The DocumentLink package also provides a indexing wrapper that will be used during indexing. This allows

the DocumentLink to be indexed as expected.

59.3. DocumentLink and Repository

In implementation projects using DocumentLink, we usually don't really care to know where the real (target) DocumentModel will be stored, because the user will probably always manipulate the DocumentModel through DocumentLinks.

In a way, it means that the real DocumentModels are "somewhere is space" and the user navigate via a hierarchy of DocumentLinks.

For that purpous, the documentLink package provide a `DocRepository` service that will manage the storage of the real DocumentModel.

"Real DocumentModel" won't be stored "in space" but in a hidden folder structure that will dispatch the DocumentModels in a hidden tree.

The repository storage is pluggable so you can define :

- the sub directory tree structure
- how rights are set

59.4. Using DocumentLink

In order to use DocumentLink feature you need to install the `nuxeo-platform-documenntlink` packages into your nuxeo. For that, just copy the jars into nuxeo.ear/plugins and restart your server.

On a stock Nuxeo EP, documentLink won't give you much visible features : it provides some new API, but since the default Web Application does not use it, it won't be really usefull. The DocumentLink Addon must be considered as a new API.

One of the DocumentLink package is dedicated to unit tests (`nuxeo-platform-documentlink-tests`), it can be used as a good sample of code for using all DocumentLink APIs.

Nevertheless, here are some simple examples :

59.4.1. Using the adapter

The `DocumentLinkAdapter` is available on any DocumentModel that has the DocumentLink Schema. Using the adapter you can define the target DocumentModel and the masked schemas, but you can also have a full access to the DocumentModel API .

```
String startPath="/default-domain/ws1";

// create a simple DocumentModel
DocumentModel doc = coreSession.createDocumentModel(startPath, "file", "File");
doc.setProperty("dublincore", "title", "MyDoc");
doc.setProperty("dublincore", "coverage", "MyDocCoverage");
doc = coreSession.createDocument(doc);

// create a DocumentModel with the DocumentLink type
DocumentModel link = coreSession.createDocumentModel(startPath, "link","DocumentLink");
link.setProperty("dublincore", "title", "MyLinkToDoc");
link = coreSession.createDocument(link);


// get the DocumentLinkAdapter
DocumentLinkAdapter adaptedLink = link.getAdapter(DocumentLinkAdapter.class);
// set the target DocumentModel
adaptedLink.setTargetDocument(doc);

// check property access

// check property accessor pass-throught
String cover1 = (String) link.getProperty("dublincore", "coverage");
```

```
String cover2 = (String) adaptedLink.getProperty("dublincore","coverage");
assertNotNull(cover2);
assertEquals("MyDocCoverage", cover2);

// check automatic override (title)
String title0 = (String) doc.getProperty("dublincore", "title");
String title1 = (String) link.getProperty("dublincore", "title");
String title2 = (String) adaptedLink.getProperty("dublincore", "title");

// adapter and documentModel direct Access return the same value
assertEquals(title1, title2);
// DocumentLink do not return the title of the target since it's overridden
assertFalse(title1.equals(title0));
```

59.4.2. Using the DocRepository

DocumentLink provides static helpers to help you create DocumentLinks and use the DocRepository.

```
DocumentModel dm = coreSession.createDocumentModel("File");
dm.setProperty("dublincore", "title", "testme");

// create a new DocumentLink of type DocumentLink
// create it in / path
// make it point to the dm DocumentModel that will be stored somewhere in the DocRepository
DocumentLinkAdapter adaptedLink = DocumentLinkHelper
                .createDocumentInCentralRepository(coreSession, dm, "/",
                                "DocumentLink");
coreSession.save();

// get the create target DocumentModel
DocumentModel targetDoc = adaptedLink.getTargetDocument();

// get the Repository
DocRepository repo = DocRepositoryHelper.getDocumentRepository(coreSession);

// get all the DocumentLinks
DocumentModelList proxies = repo.getProxiesForDocument(targetDoc.getRef());
assertTrue(proxies.size()==1);

// create a second DocumentLink pointing to the same target
DocumentLinkAdapter secondLink=DocumentLinkHelper.createDocumentLink(targetDoc, "/");
coreSession.save();

// get all the DocumentLinks
proxies = repo.getProxiesForDocument(targetDoc.getRef());
assertTrue(proxies.size()==2);
```

Chapter 60. Nuxeo Search Center

This chapter presents the architecture of the `nuxeo-searchcenter` GWT client to the Search Service (new in Nuxeo 5.2).

60.1. Overview

The Nuxeo Search Center is a standalone GWT application that can be deployed on a Webengine site in order to provide an AJAX interface to perform queries on the document repository.

TODO: add a screenshot here

60.2. Installing and testing the search center

60.2.1. Building from source

First of all, Nuxeo Search Center depends on nuxeo-gwt modules, so the first thing to do is to build and install the nuxeo-gwt projects.

```
hg clone https://hg.nuxeo.org/nuxeo/nuxeo-gwt
cd nuxeo-gwt
mvn clean install
```

Then you need the Nuxeo Search Center projects.

```
hg clone https://hg.nuxeo.org/nuxeo-searchcenter
cd nuxeo-searchcenter
mvn clean install
```

60.2.2. Deploying on a Nuxeo server

To get the Nuxeo Search Center working on Nuxeo, only those 3 jars are needed:

- `nuxeo-gwt-server-5.2-xxx.jar`
- `nuxeo-searchcenter-gwt-app-5.2-xxx.jar`
- `nuxeo-searchcenter-service-5.2-xxx.jar`

To deploy it, drop these 3 jars:

- in `nuxeo.ear/plugins` on JBoss.
- in `bundles` on Webengine standalone (Jetty).

Restart your server, and you can access Nuxeo Search Center at the following urls:

- `http://localhost:8080/nuxeo/site/searchcenter` on JBoss.
- `http://localhost:8080/searchcenter` on Webengine standalone (Jetty).

The Nuxeo Search Center application is also listed with the others Webengine applications at

`http://localhost:8080/nuxeo/site` or `http://localhost:8080/`.

60.2.3. Using the Search Center

TODO: go to `http://localhost:8080/site/` and ... .

60.2.3.1. Save/Load FilterSet

The current search can be saved from the graphical interface. It will send the current FilterSet, which contains all the FilterWidgets descriptions and their values, to the SearchCenterService. The FilterSet is saved as a DocumentModel on the user personnal Workspace (cf. nuxeo-platform-userworkspace add-on).

A ComboBox lists all the available FilterSets : the contributed ones and the saved ones which the current user can access. When a FilterSet is selected and the load button pressed, the list of FilterWidgets is rebuild, their default values are setted and the query is regenerated and executed.

60.2.3.2. TODO

TODO: Explain how to customize the list of available filter widgets (from the user interface).

60.3. How does it work

The Nuxeo Search Center application is composed of 3 different parts.

60.3.1. The GWT client

The UI part is built upon Google Web Toolkit, and particularly with SmartGWT

The client part lies in the `nuxeo-searchcenter-gwt-app` project.

60.3.2. The JAX-RS server

The JAX-RS part of Nuxeo Search Center lies in the class `SearchCenterModule` located in `nuxeo-searchcenter-gwt-app/src/main/java/org/nuxeo/ecm/searchcenter/module`.

Here are the main methods, defined in `SearchCenterModule`, used by Nuxeo Search Center:

- `getVocabulary`, used to get the values list of a given simple vocabulary. It returns the list in a JSON format understandable by the SmartGWT widgets.
- `getHierarchicalVocabulary`, used to get the values list of a given hierarchical vocabulary. It returns the list in a JSON format understandable by the SmartGWT widgets.
- `execQuery`, used to create the Query from the list of filtering values, and execute it by calling the `SearchCenterService`. It returns a document list in a JSON format understandable by the SmartGWT widgets.
- `getFilterSet`, used to get a FilterSet by its name or id. It returns a FilterSet description in a JSON format used to build the different widgets.
- `getFilterSets`, used to get all the available FilterSets. It returns a list of FilterSet descriptions in a JSON format used to fill the FilterSet load ComboBox.
- `saveFilterSet`, used to save the current search. It takes in parameter the JSON representation of the current FilterSet.

60.3.3. The SearchCenterService

The SearchCenterService is defined in the project `nuxeo-searchcenter-service`. Its purpose is to manage the FilterSets and FilterWidgets and generate and execute the QueryModel from the values entered by the user.

A helper class (`SearchCenterHelper`) is used to call the service methods and serialize the results in JSON format. This helper is called by some of the JAX-RS methods like `execQuery` and `getFilterSet`.

60.4. How to customize the list of available filters

60.4.1. FilterWidgets

To register a FilterWidget, you need to write a contribution with an extension to the SearchCenterService like :

```
<extension target="org.nuxeo.ecm.searchcenter.service.SearchCenterService" point="filterWidget">
  <filterWidget
    name="dc_subjects"
    index="dc:subjects"
    title="Subjects"
    type="hierarchicalVocabulary"
    vocabulary="topic" />
</extension>
```

- The `name` attribute specifies a unique name for the FilterWidget. This name is used in FilterSet registrations.
- The `index` attribute specifies the field in the associated where clause of the query.
- The `title` attribute specifies the displayed title in the UI.
- The `type` attribute specifies the type of the widget. The available widget types are :
  - `string`
  - `date`
  - `vocabulary`
  - `hierarchicalVocabulary`
- The `maxValues` attribute is optional and specifies the number of maximum values the user can enter.
- The `operator` attribute is optional and specifies the operator used in the corresponding where clause of the generated QueryModel. Default operator depends of the widget type :
  - `string` : if `maxValues == 1`, default operator is `=` else it is `IN`
  - `date` : default operator is `BETWEEN`
  - `vocabulary` and `hierarchicalVocabulary` : default operator is `IN`
- The `vocabulary` attribute is only used by widgets of type `vocabulary` and `hierarchicalVocabulary`. It specifies the vocabulary which the values of the widgets are fetched from.

60.4.2. FilterSets

To register a FilterSet, you need to write a contribution with an extension to the SearchCenterService like :

```
<extension target="org.nuxeo.ecm.searchcenter.service.SearchCenterService" point="filterSet">
  <filterSet name="myFilterSet" title="My FilterSet" enabled="true">
    <filterSetItem widgetName="dc_creator" />
    <filterSetItem widgetName="dc_modified" />
  </filterSet>

  <fixedPart>ecm:primaryType != 'FilterSet' AND ecm:isCheckedInVersion = 0</fixedPart>
</extension>
```

- The `name` attribute specifies a unique name for the FilterSet.
- The `title` attribute specifies a display title for the FilterSet.
- The `enabled` attribute is optional and allows to disable the FilterSet if set to `false`. Default value is `true`.
- The `filterSetItem` tags designate the list of FilterWidgets the FilterSet contains.
  - The `widgetName` attribute refers to a FilterWidget name.
  - The `hidden` optional attribute allows to hide the specified widget. Default value is `false`.
- The `fixedPart` tag can be used to add a fixed part to the generated where clause of the query.

60.5. How to build new GWT filter widgets

60.5.1. Setting up a GWT eclipse development environment

Inside the nuxeo-searchcenter directory, first run

```
mvn eclipse:eclipse
```

to setup the required files needed by eclipse to import the projects, then import the 2 projects in your workspace.

The simplest way to run and debug Nuxeo Search Center while developing is to use the GWT hosted mode.

- Deploy Nuxeo Search Center on either JBoss or Jetty, and start the server.
- inside the `nuxeo-searchcenter-gwt-app`, execute:

```
mvn gwt:eclipse
```

  This will create an Eclipse launcher, `org.nuxeo.ecm.searchcenter.gwt.SearchCenter.launch`.

- add `-noserver` to the arguments list of the new configuration.

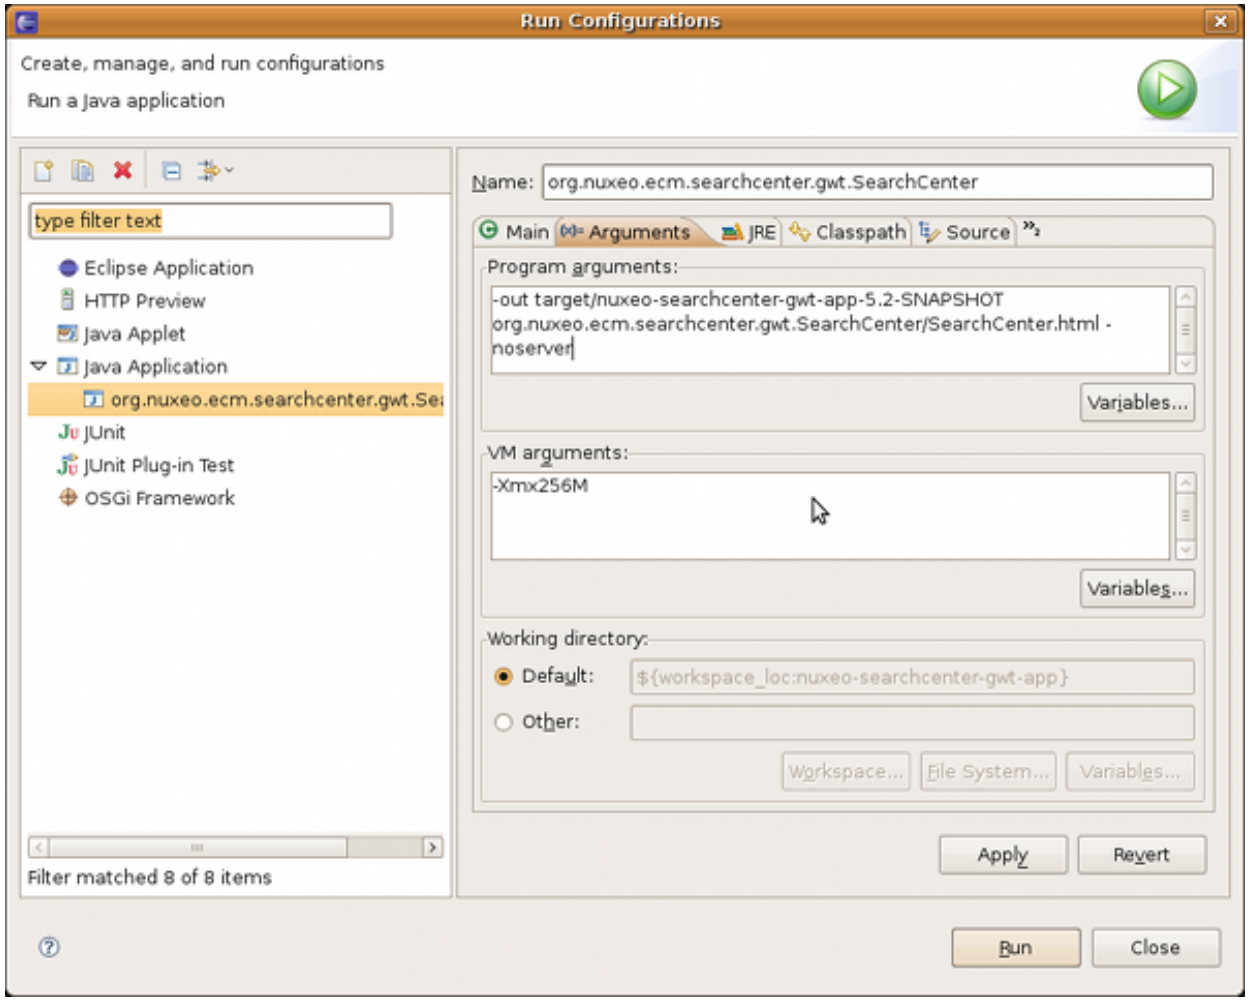

Finally, run the configuration. After the connection was refused to `http://localhost:8888`, just put the url of the deployed Nuxeo Search Center: `http://localhost:8080/nuxeo/site/searchcenter` for JBoss, or `http://localhost:8080/searchcenter` for Jetty.

In order to write your own JAX-RS components in Groovy you might want to also install the Groovy eclipse plugin from the following update site: `http://dist.codehaus.org/groovy/distributions/update/` although you can also write them in plain old Java too.

60.5.2. Existing filter widgets

There are 4 filter widgets bundled with Nuxeo Search Center:

- `StringFilterWidget`: filter documents through arbitrary text entered by the user.
- `DateFilterWidget`: filter documents via a Date interval
- `VocabularyFilterWidget`: filter documents by selecting one or more value(s) among the configured simple vocabulary values.
- `HierarchicalVocabularyFilterWidget`: filter documents by selecting one or more value(s) among the configured hierarchical vocabulary values.

Each widget instance is bound to a field in the Nuxeo documents. For instance, a String Filter Widget can be bound to the `ecm:fulltext` field, and a Date Filter Widget to the `dc:created` field.

These 4 widgets all extend `VLayoutFilterWidget` abstract class. This class is a convenient base class for the filter widgets you will develop. The `VLayoutFilterWidget` class provides:

- a defaut layout for the widget composed of a title, a display list of the selected values and a form used to add new values.

  TODO: screenshot of a simple filter widget.

- useful methods to add or remove filtering values with their associated rows in the display list.

- a method `refreshFilterSet` which updates the current `FilterSet` and asks the application to refresh the documents list.

Your new filter widget class can extend `VLayoutFilterWidget` but it is not required. The only requirement is that your class extends `FilterWidget` so that it can be used by the different `FilterView` of the application.

The widgets are all built upon JSON Objects which contain at least the following attributes: `type`, `name`, `title`, `hidden`, all setted from the registered values in the `FilterWidget` and `FilterSet` definition.

Each widget extending `VLayoutFilterWidget`can be configured to accept a maximum number of values through the `maxValues` parameter.

60.5.3. TODO: layout and extension points

TODO

Part VIII. Annexes

Appendix A. Frequently Asked Questions

A.1. Deployment and Operations

Nuxeo supports and certifies the following hardware and software:

Hardware:

- Intel/AMD 32-bit & 64-bit
- SPARC 32-bit & 64-bit

Operating Systems:

- RedHat 3.x, 4.x, 5.x
- Debian 4.0, Ubuntu Server 6.06 LTS and 7.04
- Solaris 10
- Windows Server 2003
- MacOS X 10.4.x

RDBMS:

- PostgreSQL 8.x
- MySQL 5.x
- Oracle Database 9i, Oracle Database 10g

Java Runtime Environment (JRE):

- Java 5 aka 1.5.0 (update 11 recommended)
- Java 6 aka 1.6.0 (update 11 recommended)

Java EE application servers:

- JBoss AS 4.0.4 GA and 4.0.5 GA
- JBoss AS 4.2.0 GA (in progress)
- Glassfish v2 (in progress)
- BEA WebLogic 10 (in progress)

The most used configuration is JBoss AS 4.0.4 GA using JRE 1.5.0_11 on RedHat AS 4.x running on Intel x86 hardware.

Intel-based hardware (bi-Dual Core, Quad Core or bi-Quad Core), 4GB of RAM. The required disk space only depends on the data volume to store (raw requirements to be secure: size of files to manage * 2).

Each release (major and minor) is delivered with an upgrade procedure, new features list, improvements list, fixed bugs list and known bugs/limitations list. Moreover the issue tracker is public (it allows everybody to see the status of the software, known/ongoing bugs and issues, features/improvement roadmap, etc.).

An installation guide is available. Upgrade procedures are delivered along each release.

An administration guide is available and updated with each release.

A User Guide is delivered with each release.

A developer guide is delivered and constantly improved.

The reference manual assembles all the documentation available for users, developers, operation teams, etc.

Nuxeo provides several other documents/resources such as the API (Javadoc), some tutorials, specific *Archetypes* for Maven 2 (useful to quickly bootstrap new plugins / projects), etc.

The documentation is available in the English language. Translation to French, Spanish, German and Italian are supported/provided by the community (if you require a specific language, you can order it from Nuxeo).

The documentation of included software are either included in the reference documentation if it's useful for common operations, either linked if not. All the documentation of Nuxeo EP and bundled software packages is freely available.

Yearly major release and quarterly minor release. The high-level roadmap is published by Nuxeo and updated frequently. Detailed roadmap is available from the issue tracker (all details are available on each issue such as comments, status, votes, related commit in the SCM, etc.).

See installation procedure. XXX Add link.

The "Administration and Operation Guide" describes available monitoring points. In short, Nuxeo EP offers a set of JMX services to monitor all critical points of the application (standard Java EE applications monitoring system). Moreover, logs can be broadcasted using log4j capabilities (SNMP, email, etc.). Both should be usable by all major monitoring software.

Nuxeo EP is fully based on Java EE 5 and supports related clustering and HA features. JBoss Clustering is the recommended clustering and HA solution for Nuxeo EP's services. Nuxeo EP services can be configured for performance clustering and/or HA clustering (depending on the capabilities and requirements of each service).

This is possible through configuration of the application server (ex: JBoss AS / Tomcat). Nuxeo EP relies on the application server for all the network configuration.

Nuxeo EP entirely depends on the Java EE application server for all network related configuration. It is not bind in any way to the physical network configuration of the server. Hence it is possible to change the hostname of the server and restart the machine without causing any problem to Nuxeo EP.

Fail-over relies on JBoss Clustering for Nuxeo EP services. Here is the HA system used for each category of services:

- *Nuxeo Core* (Content Repository): HA clustering only. It relies on the native RDMBS replication system (Oracle RAC or PostgreSQL replication solutions). Data integrity has to be trustable and enforced.
- *Nuxeo Search* (Search Engine): HA and performance clustering. It can use a shared filesystem (if indexes are stored on the filesystem) or can rely on the RDMBS replication solution. If data integrity is corrupt, a reindexing of the content is be sufficient to restore it.
- *EJB3-based Services*: HA and performance clustering. Use native EJB3 clustering and load-balancing from JBoss Clustering. Services using data persistence rely on RDMBS replication (for HA) that needs to be trustable and enforced.
- *Web Client/App*: can use HA and performance clustering (using JBoss Clustering). Does not need data sync.

To achieve the highest level of data integrity, Nuxeo recommends storing binary files as BLOBs directly in the database (hence use a RDBMS offering optimized BLOBs storage (such as Oracle or PostgreSQL). Using this mechanism, Nuxeo EP can store all its data into the RDBMS (including request/search engine indexes) and

relies on it to enforce data integrity. Moreover, Nuxeo EP is *fully* transactional and relies on JTA (+ XA) for transaction management (that enforce data integrity across data sources).

Nuxeo EP has been designed to completely rely on RDBMS data integrity (that can be considered trustable nowadays). One can use RDBMS tools to check data integrity/consistency and data failure if any. If the data model is corrupted, Nuxeo EP warns about it when the repository starts. Indexes (from Nuxeo Search) can be verified and easily be rebuilt by reindexing the content if any problem occurs.

The maximal impact is service downtime and data restoration from backups. Data integrity errors are reported in the logs and can then be sent via email notifications, SNMP and any log4j capabilities.

Nuxeo EP offers an applicative data import/export service (using XML serialization of documents) that can be used as an incremental backup/restore system. For an efficient backup system, Nuxeo recommends using native RDBMS tools (that can offer incremental backups, snapshots, hot restore, etc.).

The restoration speed is the native database write performances. We do not have more statistics yet (but should be available by July 2007, benchmark are in progress on this point).

When all datasources for storage are using the same database (the recommended setup), the RDBMS can achieve a consistent backup (usually at low cost for the user service). Restore can only be launched when the system is stopped.

Content object restore can be done using the import/export service. Here is the procedure to achieve this:

1. Get IDs of content object to restore using, for example, the audit service/log (ex: get all DocIds from "CreateObject" log entries for a particular user).

2. Get those document from a backup (done via the export service) and copy them in a directory (the standard export format use one directory per content object which is easing a lot this operation).

3. Use a command line import/export client to (re-)import document in this directory.

4. You're done.

RDBMS backup can be handled as usual using legacy backup scripts for this RDBMS. Applicative backups can be launched using the import/export client CLI. There is not supported scripts at the moment (but they could easily be written).

Appendix B. Detailed Development Software Installation Instructions

This chapter is provided to help you install the software needed to work on your Nuxeo projects.

[TODO: refactor this chapter as in many cases, the packages may be installed using some package management system (for instance on Mac OS: **port install apache-ant maven2**).]

B.1. Installing Java 5

Nuxeo EP uses the latest generation of Java technologies and thus requires a Java 5 VM such as the reference implementation by Sun.

You may already have the right Java development kit install can enter in the command line:

```
javac -version
```

. This should return something like:

```
javac "1.5.0_10"
```

If the java version is 1.5.x, you can skip the "Installing java 5" section.

Warning

Nuxeo EP doesn't currently compile under Java 6 (the 1.6.0 JDK from Sun). Even if it did, it is not clear that JBoss, the application server we are targeting, works under Java 6.

Note

Java 5 is also sometimes called Java 1.5 or 1.5.0.

B.1.1. Using the Sun Java Development Kit (Windows and linux)

Sun Microsystems provides freely downloadable version of the Java Development Kit (JDK), that is needed to compile the Nuxeo platform.

For the purpose of Nuxeo development, you should download the latest release of the JDK 5.0 (JDK 5.0 Update 11 at the time of this writing) from `http://java.sun.com/javase/downloads/index_jdk5.jsp`

B.1.2. Using a package management systems (Linux)

Some Linux distributions now include Java 5 in their official repositories.

For instance with Ubuntu (Edgy and later):

- enable "multiverse" (2 lines to uncomment) in your `/etc/apt/sources.list`
- update your package indexes:

```
$ sudo apt-get update
```

- install the full Java 5 stack (probably not all is necessary):

```
$ sudo apt-get install "sun-java5-*"
```

- ensure Java 5 is now the default JVM on your system (instead of gcj and friends by default):

```
$ sudo update-alternatives --set java /usr/lib/jvm/java-1.5.0-sun/jre/bin/java
```

(TODO: add similar instructions for Fedora Core and Debian)

B.1.3. Manual installation (Linux)

You can also manually install Java by following the instructions of this page:
http://www.java.com/en/download/linux_manual.jsp

B.1.4. Setting up JAVA_HOME (Windows, Linux, Mac OS)

This will be required by tools such as Maven (see later).

B.1.4.1. Windows

Follow these instructions:

- type 'windows' +'pause'
- click on advanced tab
- click on "environment variables" (at the bottom)
- click on new and enter "`JAVA_HOME`" for variable name and "`C:\Program Files\Java\jdk1.5.0_10`" (adapt to your own JDK install)
- don't forget to click on ok to close the environment variables window.

B.1.4.2. Linux

In your `.bashrc` (or `.zshrc`, ...) add something like:

```
export JAVA_HOME=/usr/lib/jvm/java-1.5.0-sun
```

B.1.4.3. Mac OS

Under Mac OS X, if you have properly installed a JDK (XXX: check how), you will need to put in your your `.bashrc` (or `.zshrc`, ...) add something like:

```
export JAVA_HOME=/Library/Java/Home
```

B.2. Installing Ant

Ant will be used for building process. If you didn't set it up already on your computer, you can download it here.

Then need to have Ant setup and on your PATH environment variable.

For linux:

Add something like the following in your `.bashrc`:

```
export PATH=/opt/apache-ant-1.7.1/bin:$PATH
```

For Windows:

- type 'windows' +'pause'
- click on advanced tab
- click on "environment variables" (at the bottom)
- click on path variable and click modify
- add something like this at the end of the PATH definition: `;c:\program files\apache-ant-1.7.1\bin`
- don't forget to click on OK to close the environment variables window.

You can then check that your installation is correct by typing:

```
ant -version
```

B.3. Installing Maven

B.3.1. What is Maven?

Quoting the Wikipedia entry for Maven:

> Maven is a software tool for Java programming language project management and automated software build. It is similar in functionality to the Apache Ant tool, but has a simpler build configuration model, based on an XML format. Maven is hosted by the Apache Software Foundation, where it was formerly part of the Jakarta Project.
>
> Maven uses a construct known as a Project Object Model (POM) to describe the software project being built, its dependencies on other external modules and components, and the build order. It comes with pre-defined targets for performing certain well defined tasks such as compilation of code and its packaging.
>
> A key feature of Maven is that it is network-ready. The core engine can dynamically download plugins from a repository, the same repository that provides access to many versions of different Open Source Java projects, from Apache and other organizations and developers. This repository and its reorganized successor the Maven 2 repository are the de facto distribution mechanism for Java applications. Maven provides built in support not just for retrieving files from this repository, but to upload artifacts at the end of the build. A local cache of downloaded artifacts acts as the primary means of synchronizing the output of projects on a local system.
>
> Nuxeo is now fully "maven managed".

Nuxeo holds a Maven repository here: `http://maven.nuxeo.org/`.

B.3.2. Installing Maven

You should then install Maven 2 on your development box by downloading the latest tarball from http://maven.apache.org/download.html and then untar the archive in `/opt` (for instance).

We recommend that you use the latest version of Maven (2.0.9 at the time of this writing).

As usual, you have to put the `mvn` executable into the path of your environment (cf. Ant)

Then add the `bin/` subdir in your `PATH` by adding something like the following in your `.bashrc`:

```
export PATH=/opt/maven-2.0.9/bin:$PATH
```

In a new shell you can then test:

```
$ mvn --version
Maven version: 2.0.9
```

B.3.3. Generate a new project with the nuxeo-archetype-start archetype

The goal of the nuxeo-archetype-start template is to setup a development environment to work on a Nuxeo EP plugin.

The default code provides: a maven layout for sources, tests and dependencies, a Ant target for deployment. It also customizes the web application a litte bit.

To create a project named `my-project` in the `com.company.sandbox.myproject` package:

```
mvn
 org.apache.maven.plugins:maven-archetype-plugin:1.0-alpha-7:create \
 -DartifactId=my-project \
 -DgroupId=com.company.sandbox.myproject \
 -DarchetypeArtifactId=nuxeo-archetype-start \
 -DarchetypeGroupId=org.nuxeo.archetypes \
 -DarchetypeVersion=5.1.6 \
 -DremoteRepositories=http://maven.nuxeo.org/nuxeo-release
```

You can see that you need to supply six arguments:

- *artifactId* : the name of your project, usually with '-' to separate the words if there are many.
- *groupId* : the domain name of your project. Usually it is the package parent name of your classes, you should use '.' to separate the words ('-' is not supported here).
- *archetypeArtifactId* : the maven archetype artifact id. To generate a new project, Nuxeo provides you with nuxeo-archetype-start
- *archetypeGroupId* : unique for all nuxeo maven archetypes : org.nuxeo.archetypes
- *archetypeVersion* : the version of the archetype which is equivalent to the version of Nuxeo EP without the *GA* or *RC* part. (5.1.6, 5.1.5 ...).
- *remoteRepositories* : the repository location to download the archetype.

B.4. Installing JBoss AS

Nuxeo EP default target is the JBoss application server with EJB3 support enabled. To enable the EJB3 support, you should install JBoss with the latest version of the JEMS installer:

```
$ sudo java -jar jems-installer-1.2.0.GA.jar
```

While executing the installation wizard, you must select `ejb3` install. You can leave all other parameters to their default values.

You would get PermGenSpace errors if you run JBoss without this configuration:

- Linux:

  Edit `/opt/jboss/bin/run.conf` and add the following line at the end of the file

```
JAVA_OPTS="$JAVA_OPTS -XX:MaxPermSize=128m"
```

  Restart JBoss.

- Windows:

  Edit `$JBOSS/bin/run.bat` and add the following line after the line `: set JAVA_OPTS=%JAVA_OPTS% -Dprogram.name=%PROGNAME%`

```
set JAVA_OPTS=%JAVA_OPTS% -XX:MaxPermSize=128m
```

  Restart JBoss.

B.4.1. JBoss AS listening ports customization

The common task for JBoss users is making it to communicate over a single HTTP server. This is quite useful for network administration, making it easier to go through firewalls. This section describes the necessary steps to make JBoss communicate primarily over HTTP

B.4.1.1. Tomcat Web server

JBoss is shipped with built-in Tomcat web server. This server is configured in 'deploy/jbossweb-tomcat55.sar/server.xml' By default only two connectors are enabled: HTTP connector (port 8080) and AJP connector (port 8009). Generally speaking you need only one of them. The former connector is needed if standalone HTTP server built in JBoss is used. You may want to configure it to listen the default HTTP port 80. The latter connector is needed only if you want to couple JBoss server with external web server like Apache, in this case it is reasonable, for security issues to change the binding address to 'localhost' (of course if Apache runs on the same host).

B.4.1.2. HTTP invoker web application

The JBoss default configuration deploys a special service that can be used to expose different JBoss services in the HTTP server. It is located in 'deploy/http-invoker.sar'. The configuration file 'deploy/http-invoker.sar/META-INF/jboss-service.xml' may be tweaked to tune the AS to specific needs. By default the service provides HTTP invoker MBean for EJB ('jboss:service=invoker,type=http') and two HTTP proxy MBeans that marshal the requests to the Naming service MBean ('jboss:service=invoker,type=http,target=Naming' and 'jboss:service=invoker,type=http,target=Naming,readonly=true'). If you need to provide HTTP interface to another MBeans, you also may specify the proxy services in 'deploy\http-invoker.sar\META-INF\jboss-service.xml'. For instance the SRP service for JBoss authentication may be exposed here.

The service also deploys web application 'deploy/http-invoker.sar/invoker.war', that configures the servlets that convert HTTP requests into invocation of MBeans/EJB methods. If you add HTTP proxies to MBeans, you may need to add servlets that handle the corresponding URI.

Important note. If HTTPS protocol is used the configuration should not use the default host name because the virtual host name used in the URL (say 'www.nuxeo.org') and exposed in SSL certificates usually differs from the computer name where JBoss is running. To accomplish this get rid of the following attributes: InvokerURLPrefix, InvokerURLSuffix, UseHostName, replacing them with a single InvokerURL attribute, like this:

```
<mbean code="org.jboss.invocation.http.server.HttpProxyFactory"
   name="jboss:service=invoker,type=https,target=Naming">
```

```
   <!-- Compose the invoker URL from the cluster node address -->
   <attribute name="InvokerURL">
      https://www.nuxeo.org/invoker/JMXInvokerServlet
   </attribute>
   <attribute name="ExportedInterface">
      org.jnp.interfaces.Naming
   </attribute>
   <attribute name="JndiName"></attribute>
   <attribute name="ClientInterceptors">
   <interceptors>
      <interceptor>org.jboss.proxy.ClientMethodInterceptor</interceptor>
      <interceptor>org.jboss.proxy.SecurityInterceptor</interceptor>
      <interceptor>org.jboss.naming.interceptors.ExceptionInterceptor</interceptor>
      <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
   </interceptors>
   </attribute>
   <depends>jboss:service=invoker,type=https</depends>
</mbean>
<!-- The rest MBeans should also use InvokerURL attribute only,
   make sure you specify the right host name
-->
```

B.4.1.3. JNDI service

This is the core service of JBoss and should never be disabled. Nevertheless this service does not need own listening port (1099,1098), so just change the '1099' to '-1':

```
<mbean code="org.jboss.naming.NamingService"
   name="jboss:service=Naming"
   xmbean-dd="resource:xmdesc/NamingService-xmbean.xml">
   <!-- The call by value mode. true if all lookups are unmarshalled using
   the caller's TCL, false if in VM lookups return the value by reference.
   -->
   <attribute name="CallByValue">false</attribute>
   <!-- The listening port for the bootstrap JNP service. Set this to -1
     to run the NamingService without the JNP invoker listening port.
   -->
   <attribute name="Port">-1</attribute>
   <!-- The bootstrap JNP server bind address. This also sets the default
   RMI service bind address. Empty == all addresses, use localhost to hide this from
   network
    -->
   <attribute name="BindAddress">localhost</attribute>
   <!-- The port of the RMI naming service, 0 == anonymous, you cannot use -1 here -->
   <attribute name="RmiPort">1098</attribute>
   <!-- The RMI service bind address. Empty == all addresses, use localhost to hide this from
   network
    -->
   <attribute name="RmiBindAddress">localhost</attribute>
   <!-- The thread pool service used to control the bootstrap lookups -->
   <depends optional-attribute-name="LookupPool"
      proxy-type="attribute">jboss.system:service=ThreadPool</depends>
</mbean>
```

B.4.1.4. Default JBoss EJB invokers

You may deinstall the JRMP and Pooled invokers completely. Just comment out the MBeans that provide the corresponding services in 'conf/jboss-service.xml'.

Important note. The JBoss specifies the invokers for EJB in 'conf/standardjboss.xml' file. The default is 'jboss:service=invoker,type=jrmp' invoker. To change it to HTTP invoker you need to add invoker bindings for all EJB types deployed in your applications. Generally it means you need to copy all "*-rmi-invoker" bindings into "*-http-invoker" bindings, replacing "<invoker-mbean>jboss:service=invoker,type=jrmp</invoker-mbean>" with "<invoker-mbean>jboss:service=invoker,type=http</invoker-mbean>" for the new bindings. Also you will need to make the HTTP invoker default for all EJB container configurations replacing "<invoker-proxy-binding-name>*-rmi-invoker</invoker-proxy-binding-name>" with "<invoker-proxy-binding-name>*-http-invoker</invoker-proxy-binding-name>" correspondingly.

The easiest (but probably not the right) way for JBoss 4.0.x is to replace the string 'jboss:service=invoker,type=jrmp' with 'jboss:service=invoker,type=http' in this file by any text editor. It may be not correct if you want to mix both invokers for your EJBs.

B.4.1.5. JBoss EJB3 invoker

The EJB3 invoker which is specified in 'deploy/ejb3.deployer/META-INF/jboss-service.xml' uses JBoss remoting mechanism. By default it is bound to socket with a connector listening TCP/IP port 3873. This should be changed to the servlet locator:

```
<mbean code="org.jboss.remoting.transport.Connector"
       name="jboss.remoting:type=Connector,name=DefaultEjb3Connector,handler=ejb3">
   <depends>jboss.aop:service=AspectDeployer</depends>
   <attribute name="InvokerLocator">
      servlet://${jboss.bind.address}/invoker/Ejb3InvokerServlet
   </attribute>
   <attribute name="Configuration">
      <handlers>
         <handler subsystem="AOP">org.jboss.aspects.remoting.AOPRemotingInvocationHandler</handler>
      </handlers>
   </attribute>
</mbean>
```

The corresponding servlet should be added to invoker web application descriptor ('http-invoker.sar/invoker.war/WEB-INF/web.xml'):

```
<servlet>
   <servlet-name>Ejb3InvokerServlet</servlet-name>
   <description>
      The ServerInvokerServlet receives requests via HTTP protocol
      from within a web container and passes it onto the
      ServletServerInvoker for processing.
   </description>
   <servlet-class>
      org.jboss.remoting.transport.servlet.web.ServerInvokerServlet
   </servlet-class>
   <init-param>
      <param-name>locatorUrl</param-name>
      <param-value>
         servlet://${jboss.bind.address}/invoker/Ejb3InvokerServlet
      </param-value>
      <description>
         The servlet server invoker locator url
      </description>
   </init-param>
   <load-on-startup>1</load-on-startup>
</servlet>
<servlet-mapping>
   <servlet-name>Ejb3InvokerServlet</servlet-name>
   <url-pattern>/Ejb3InvokerServlet/*</url-pattern>
</servlet-mapping>
```

B.4.2. Affected JBoss services

As JRMP invoker is used in many other JBoss services, so it should be replaced. The affected MBeans are "jboss:service=ClientUserTransaction" (conf/jboss-service.xml), "jboss.jmx:type=adaptor,name=Invoker,protocol=jrmp,service=proxyFactory".

B.4.2.1. Client User Transaction

The Client User Transaction service depends on two JRMP Proxy Factories described in the nested MBeans. Every JRMP proxy factory should be replaced with HTTP proxy factory:

```
<mbean
   code="org.jboss.tm.usertx.server.ClientUserTransactionService"
   name="jboss:service=ClientUserTransaction"
   xmbean-dd="resource:xmdesc/ClientUserTransaction-xmbean.xml">
  <depends>
    <mbean code="org.jboss.invocation.http.server.HttpProxyFactory"
         name="jboss:service=proxyFactory,target=ClientUserTransactionFactory">
      <attribute name="InvokerName">jboss:service=invoker,type=http</attribute>
      <attribute name="JndiName">UserTransactionSessionFactory</attribute>
      <attribute name="ExportedInterface">
        org.jboss.tm.usertx.interfaces.UserTransactionSessionFactory
      </attribute>
      <attribute name="ClientInterceptors">
        <interceptors>
          <interceptor>org.jboss.proxy.ClientMethodInterceptor</interceptor>
          <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
        </interceptors>
      </attribute>
      <depends>jboss:service=invoker,type=http</depends>
    </mbean>
  </depends>
  <depends optional-attribute-name="TxProxyName">
```

```
    <mbean code="org.jboss.invocation.http.server.HttpProxyFactory"
         name="jboss:service=proxyFactory,target=ClientUserTransaction">
      <attribute name="InvokerName">jboss:service=invoker,type=http</attribute>
      <attribute name="JndiName"></attribute>
      <attribute name="ExportedInterface">
        org.jboss.tm.usertx.interfaces.UserTransactionSession
      </attribute>
      <attribute name="ClientInterceptors">
        <interceptors>
          <interceptor>org.jboss.proxy.ClientMethodInterceptor</interceptor>
          <interceptor>org.jboss.invocation.InvokerInterceptor</interceptor>
        </interceptors>
      </attribute>
      <depends>jboss:service=invoker,type=http</depends>
    </mbean>
  </depends>
</mbean>
```

Note that JRMP Proxy factory attributes differ from attributes of HTTP proxy factory.

B.4.2.2. JMX adaptor

The JMX adaptor is adapted for HTTP as following ('deploy/jmx-invoker-service.xml'):

```
<mbean code="org.jboss.invocation.http.server.HttpProxyFactory"
       name="jboss.jmx:type=adaptor,name=Invoker,protocol=http,service=proxyFactory">
  <attribute name="InvokerURL">https://www.nuxeo.org/invoker/JMXInvokerServlet</attribute>
  <depends optional-attribute-name="InvokerName">jboss.jmx:type=adaptor,name=Invoker</depends>
  <attribute name="ExportedInterface">org.jboss.jmx.adaptor.rmi.RMIAdaptor</attribute>
  <attribute name="JndiName">jmx/invoker/HttpAdaptor</attribute>
</mbean>
```

and ('deploy/console-mgr.sar/META-INF/jboss-service.xml'):

```
<mbean code="org.jboss.console.manager.PluginManager"
       name="jboss.admin:service=PluginManager">
  <depends>jboss.jmx:type=adaptor,name=Invoker,protocol=http,service=proxyFactory</depends>
  <!-- the rest stays intact -->
</mbean>
```

B.4.2.3. Datasource adaptors

You need to set the invoker explicitly for all deployed data sources. The element <jmx-invoker-name> should be added to all <local-tx-datasource> and <xa-datasource> elements. Otherwise the server will complain about missing JRMP invoker which is used by default:

```
<datasources>
        ...
   <local-tx-datasource>
      <!-- specify explicitly the invoker to use -->
      <jmx-invoker-name>jboss:service=invoker,type=https</jmx-invoker-name>
      <!-- the rest stays intact -->
              ...
   </local-tx-datasource>
        ...
</datasources>
```

B.5. Installing a Subversion client

The official svnbook is a very good reference for both beginners and advanced subversion users.

B.5.1. Generic subversion clients with linux

The Nuxeo EP source repository is a Subversion repository thus you will need a subversion client to access the source code. Most Linux distributions provide the `svn` command line tool. To install it under Ubuntu / Debian:

```
$ sudo apt-get install subversion
```

Under Fedora Core, this should become:

```
$ yum install subversion
```

B.5.2. Windows

For MS Windows, we recommend to use the TortoiseSVN Subversion client. You can also directly use the Subversion command-line client from Subversion website.

B.6. Installing mercurial

Nuxeo EP source are tracked using Mercurial from Selenic.

B.6.1. Linux

To install the `hg` command under Ubuntu / Debian :

```
$ sudo apt-get install mercurial
```

Under Fedora Core, this should become:

```
$ yum install mercurial
```

B.6.2. Windows

For MS Windows, we recommend to use the all in one tortoise bundle provided to you by selenic.

B.6.3. Mac OS

For Mac OS, our preferred method is to use the darwin ports environment.

```
$ port install mercurial
```

B.6.4. Setting a username

If you plan to checkin in nuxeo's hg repositories, you should provide a valid user name. This is achieved by setting the `username` property in the `.hgrc` settings file.

```
...
[ui]
 ...
 username = firstname lastname <you@your-domain>
 ...
```

B.6.5. Activating pre-integrated extensions

You should activate some pre-integrated extensions for working with nuxeo repositories. This is achieved by adding the following lines to the extensions section of your `.hgrc`.

```
...
[extensions]
  ...
  hgext.mq =
```

```
    hgext.parentrevspec =
    hgext.graphlog =
    patchbomb =
    transplant =
    ...
```

B.6.6. Using forest

Nuxeo EP sources are available as a forest (nuxeo-ecm: nuxeo-common, nuxeo-core, ...). You're able to extend the `hg` internal commands by installing the hgforest plugin. One way is to clone the forest hg plugin's repository.

```
$ hg clone http://hg.akoha.org/hgforest
```

. The next step is to declare the plugin into you `.hgrc` file in the `extension` section.

```
[extensions]
...
hgext.forest = [your installation path]/forest.py
...
```

B.7. Chapter Key Point

In this chapter, we learned:

- how to install a Java development environment (JDK) on your machine
- how to install Ant and Maven, two mandatory tools for building and deploying your own projects on top of the Nuxeo platform
- how to install the JBoss AS 4 application server that will act as a container for the Nuxeo application
- how to install Mercurial and Subversion clients

Appendix C. Commercial Support

C.1. About Us

The Nuxeo open source ECM platform is developed and published by Nuxeo SAS, a company with offices in France and the UK.

Through our commercial offer, *Nuxeo Connect*, we deliver enterprise-grade functional and technical support, certified software patches and updates, and management tools that assist you during every stage of the application life-cycle - from design and development, throughout testing and deployment, to operations and monitoring. Nuxeo Connect helps reduce business and technical risks, increase productivity, speed time to deployment and improve your success rate for all Nuxeo-based projects.

We're also happy to work with partners, IT Integrators or ISVs, to deliver the best possible applications to customers.

C.2. Contact information

C.2.1. General

Web: `www.nuxeo.com`

E-mail: contact@nuxeo.com

C.2.2. France

18-20, rue Soleillet

75020 Paris - France

Tel: +33 1 40 33 79 87

Fax: +33 1 43 58 14 15

C.2.3. UK

Garden Studios, 11-15 Betterton Street, London WC2H 9BP

Tel: +44-207-043-7933